

Visual Analytics for Communication Analysis

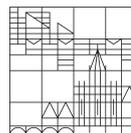
**Doctoral thesis for obtaining the
academic degree
Doctor of Natural Sciences (Dr. rer. nat.)**

submitted by

Maximilian Thomas Arnold Fischer

at the

Universität
Konstanz



Faculty of Sciences

Department of Computer and Information Science

Konstanz, 2023

Doctoral Thesis of the University of Konstanz
Day of the oral examination: December 18th, 2023

1. Referee: Prof. Dr. Daniel A. Keim, Universität Konstanz
2. Referee: Prof. Dr. Marcel Worring, University of Amsterdam
3. Referee: Prof. Dr. Sabine Storandt, Universität Konstanz

(This page intentionally left blank)

Fischer, Maximilian T.

Visual Analytics for Communication Analysis

Dissertation, University of Konstanz, 2023

Abstract

Automated systems for *analyzing digital communication*, often reliant on AI-driven models and deployed in operational analytics solutions, have become increasingly common in investigative journalism, the intelligence community, and business environments. These automated solutions, however, typically focus on specific aspects like content or network structure in isolation. Additionally, they face challenges related to ethical and privacy concerns and difficulties in efficiently combining heterogeneous data sources for *cross-methodological, holistic, and multimodal analytics*. *Visual analytics*, which combines machine learning methods with interactive visual interfaces to enable human sense- and decision-making, can be vital for designing and operating meaningful communication analysis techniques that tackle these challenges. This dissertation explores how communication analysis can be conducted in the digital age, thereby advancing the field of semi-interactive, holistic communication analysis and is divided into four parts: Part I lays the *conceptual foundations*, reviews the state-of-the-art, and formalizes the field while identifying key challenges. Additionally, it discusses *ethical and privacy* considerations and proposes ways in which visual analytics can address those issues. Part II introduces techniques related to the *identification and interpretation* of communication, including a survey of hypergraph visualizations, a technique for the identification of communication, and a method for metadata pattern analysis. Informed by these insights, Part III explores *holistic approaches* to communication analysis, first from a text-centric perspective and then from a multimodal viewpoint. The concluding Part IV summarizes the findings and anticipates future research directions. The presented techniques and methods are rigorously evaluated through comparative analyses, case studies, and expert evaluations, further discussing their applicability and generalizability. In conclusion, this dissertation frames communication in a *coherent context* of digital and semi-automated analysis within computer science, structures the field, probes ethical and privacy aspects, and describes how AI-based, multimodal, human-in-the-loop approaches can enhance intelligence analytics.

Zusammenfassung

Automatisierte Systeme zur *Analyse digitaler Kommunikation*, häufig basierend auf KI-gestützten Modellen und als Teil operativer Analyselösungen, sind in Feldern wie investigativem Journalismus, Nachrichtendiensten und Geschäftsumgebungen immer häufiger anzutreffen. Diese automatisierten Lösungen konzentrieren sich jedoch in der Regel ausschließlich auf bestimmte Aspekte wie Inhalte oder Netzwerkstrukturen. Darüber hinaus stehen sie vor ethische und datenschutzrechtliche Herausforderungen und Schwierigkeiten bei der Kombination heterogener Datenquellen für eine *methodenübergreifende, holistische und multimodale Analyse*. *Visuelle Analytik*, die Methoden des maschinellen Lernens mit interaktiven visuellen Schnittstellen kombiniert um menschliche Sinnes- und Entscheidungsfindung zu ermöglichen, kann für die Entwicklung und den Betrieb sinnvoller Kommunikationsanalysetechniken, die diese Herausforderungen lösen, entscheidend sein. Diese Doktorarbeit untersucht, wie Kommunikationsanalyse im digitalen Zeitalter funktionieren kann, um so das Feld der semi-interaktiven, ganzheitlichen Kommunikationsanalyse voranzutreiben und gliedert sich in vier Teile: Teil I legt die *konzeptionellen Grundlagen*, gibt einen Überblick über den Stand der Technik, formalisiert das Feld und identifiziert zentrale Herausforderungen. Darüber hinaus werden *ethische und datenschutzrechtliche* Erwägungen erörtert und Möglichkeiten vorgeschlagen, wie die visuelle Analytik diese Probleme angehen kann. Teil II stellt Techniken vor, die sich auf die *Identifizierung und Interpretation* von Kommunikation beziehen, einschließlich einer Übersicht über Visualisierungen von Hypergraphenmodellen, einer Technik zur Identifizierung von Kommunikation und einer Methode zur Analyse von Metadatenmustern. Auf der Grundlage dieser Erkenntnisse werden in Teil III *holistische Ansätze* zur Kommunikationsanalyse untersucht, zunächst aus einer textzentrierten Perspektive und dann aus einer multimodalen Sichtweise. Der abschließende Teil IV fasst die Ergebnisse zusammen und gibt einen Ausblick auf zukünftige Forschungsfragen. Die vorgestellten Techniken und Methoden werden anhand von vergleichenden Analysen, Fallstudien und Experteneinschätzungen rigoros evaluiert, um ihre Anwendbarkeit und Verallgemeinerbarkeit zu diskutieren. Zusammenfassend stellt diese Dissertation Kommunikation in einen *kohärenten Kontext* der digitalen und halbautomatischen Analyse innerhalb der Informatik, strukturiert die Thematik, untersucht ethische und datenschutzrechtliche Aspekte und beschreibt, wie KI-basierte, multimodale, Human-in-the-Loop-Ansätze die Kommunikationsanalyse verbessern können.

Preface

When I started my Ph.D. in late 2019 at Daniel's lab, the world seemed a different place. So much has happened since then, and the consequences of these events will likely come to define our generation and forever alter the world we live in. Some ramifications will be predictable; some will happen in the shadows and only be apparent in hindsight, while others will be unexpected in ways we cannot possibly envision nowadays. Old world orders are breaking down, and things we took for granted have disappeared, only to be replaced by a new, much more uncertain, multi-lateral world. Few would have imagined back then that—in our modern, global, and interconnected world—we are currently facing the most significant land war in Europe since World War II on Ukrainian soil. Or how fundamentally our way of living was impacted and changed due to the COVID-19 pandemic, especially in 2020 and 2021. Simultaneously, we continue to face global challenges, which we have known about for a long time but still could be more ambitious about, like gender equality and diversity, sustainability, and climate change, to name only a few. Computer science, in particular data science, can support us in these endeavors, and I'm happy to be a part of it. Technological progress and disruptive innovations are accelerating, and the progress that has been made in just slightly more than three years from when I started my Ph.D. continues to amaze me. Especially in the last months of my doctorate (early 2023), we have seen progress and unmatched speed in AI that we did not dare to imagine only two or three years ago.

Personally, for me, it has been one of the most exciting and educational but also the most challenging phases of my life. It has been an incredibly enriching opportunity to participate in state-of-the-art scientific research. My doctorate has primarily coincided with COVID-19, so I spent large parts of 2020 and 2021 working from home, researching, teaching, and presenting remotely. I spent much of this time in Hamburg, my home of choice, and despite COVID-19, this has been one of the happiest times of my life, not least due to very particular personal circumstances. However, as we have learned, the things and persons we took for granted might not be there forever, and a world in turmoil can take much: I'm incredibly grateful for everything and could not have undertaken this journey without my friends during my times in Zurich, Konstanz, and Hamburg.

I'm happy that, also during this time, I was finally able to give back to the Studienstiftung des Deutschen Volkes, from which I received a scholarship as a student, by teaching the next generation about deep learning and AI as a working group leader in an academy in 2021 (Roggenburg) and later this summer again at a summer academy in Ljubljana 2023, as well as being part of the selection committee. I feel honored about the trust my advisor Daniel A. Keim placed in me for overseeing and being responsible for several large research projects and facilitating successful

grant applications. The names PRIMAGE, AIDA-TI, RPM-BW, EACDTA, DAYDREAMS, ASGARD, and PEGASUS, will always be kept in fond memory, especially as the latter two form the conceptual basis for my research. Throughout my work, I had the unique opportunity to travel to many different places. While I could only be virtually present in Lyon (France), Salt Lake City (USA), and New Orleans (USA), I'm more than happy that I could participate in person in many different places. In Germany, these were Berlin, Bochum, Cologne, Frankfurt, Hamburg, Stuttgart, Leipzig, Munich, Oldenburg, Würzburg and internationally in Athens (Greece), Lisbon (Portugal), Porto (Portugal), Nice (France), The Hague (The Netherlands), Dublin (Ireland), Seoul (South Korea), Oklahoma City (USA), Vienna (Austria), Zurich (Switzerland), Valencia (Spain), Brussels (Belgium), Ljubljana (Slovenia), as well as Boston/Cambridge (USA).

I'm especially thrilled that one of my last large conference participations led me to Boston, my former hometown. For me, working at the intersection of AI, human intelligence, and national security in a geopolitical context, the German American Conference 2022 at Harvard Kennedy School felt like a great closure and also a good starting point. I enjoyed the high-profile conference not only from an academic perspective with core topics in AI, security, digitalization, foreign policy, and climate change - as ultimately, it is all about the people! It was an excellent opportunity to bring together minds from both sides of the Atlantic, policy leaders and pioneers in their fields, with many bright young students and young professionals to foster exchange across generations and arrive at new insights to tackle the challenges of our time, which we are facing. As such, I'm very happy that during my Ph.D., Daniel allowed me to not only work on my primary research topic (communication analysis) but also follow my passion, work as a data science consultant and project manager on the side (in my own time) to solve and communicate challenging problems, and actively shape and be engaged in projects in the infrastructure domain, in particular smart mobility, renewables, and the energy transition. There—as well as during all my work travels and over the whole course of my research—I have met many extraordinary personalities, and our discussions and their guidance, mentorship, and friendship have shaped this work—as well as me personally.

Konstanz, Germany
July 2023

Maximilian T. Fischer

Acknowledgments

I am writing these lines with sincere and profound gratitude for the inspiring discussions, fruitful collaborations, and generous support of numerous individuals and institutions, without whom this dissertation would not have been possible.

First and foremost, I would like to extend my deepest gratitude to my mentors, Daniel A. Keim and Marcel Worring. Thank you, Daniel, for the chance to pursue my doctorate in your group, for the freedom in choosing my research field, for the trust placed in me, for the inspiration and far-sighted advice, as well as the helpful discussion, and for always having an open ear for me. Also, thanks to you, Marcel, for your advice, guidance, valuable ideas, and constructive feedback. I also genuinely thank Bastian Goldlücke for serving as the chair of my doctoral committee and Sabine Storandt for being the third referee for my dissertation.

Numerous organizations directly or indirectly contributed to my academic career and my research. I sincerely thank the Studienstiftung des Deutschen Volkes (German Academic Scholarship Foundation) for the extensive support and the scholarship I received as a student. The foundation provided a thriving environment with many academic and social programmes through which I met extraordinary people, many of whom became dear friends. Further, I want to thank the European Commission (EC) and the German Federal Ministry of Education and Research (BMBF) for their research funding to our lab as part of Horizon 2020 and national programmes, which enable me to pursue my Ph.D. and research much more freely.

During my studies, I had the honor to learn from and exchange ideas with several remarkable individuals, both here at the university and during my travels. I cannot begin to list them all, as this list would never be truly complete. Nevertheless, I want to mention by name a few individuals who I feel shaped my research direction and academic career in the early stages here at the computer and information science department, inspired me, and which I want to thank: Christian Borgelt, Martin Brunner, Oliver Deussen, Tiare Feuchtner, Bastian Goldlücke, Michael Grossniklaus, Heiko Hamann, Sven Kosub, Stefan Leue, Barbara Pampel, Harald Reiterer, Dietmar Saupe, Marc Scholl, Falk Schreiber, Andreas Spitz, Sabine Storandt, Tobias Sutter, and Matthias Rupp.

Fortunately, I was also blessed to have a fantastic group of colleagues at the Chair for Data Analysis and Visualization, and extraordinary co-authors, many of whom have become good friends by now. Even while working remotely, we kept in touch constantly via all means of modern communication. Thanks to all of them: Daniel Seebacher, who supported me with my first steps in the academic world. Mennatallah El-Assady, who sharpened my academic vision and paper writing skills. Wolfgang Jentner, who taught me project and research management. And, of course, all the others: Devanshu Arya, Michael Blumenschein, Juri Buchmüller, Raphael

Buchmüller, Eren Cakmak, Frederik Dennig, Evanthia Dimara, Alexander Frings, Johannes Fuchs, Hanna Hauptmann (Schäfer), Paula Helm, Simon David Hirsbrunner, Lucas Joos, Matthias Kraus, Sabine Kuhr, Jakub Lokoč, Philipp Meschenmoser, Yannick Metz, Matthias Miller, Tom Polk, Noëlle T. Rakotondravony, Giovanna Ratini, Julius Rauscher, Dominik Sacha, Udo Schlegel, Rita Sevastjanova, Matthew Sharinghousen, Fabian Sperrle, Thilo Spinner, Manuel Stein, Dirk Streeb, and Hendrik Strobelt.

At last, I want to thank all my friends and family for their love, support, and encouragement. You made it possible to pursue my dreams, supported me no matter what, and helped me to become the person I am today.

Contents

Abstract / Zusammenfassung	i
Preface	v
Acknowledgments	ix
Contents	xii
1 Introduction to Communication Analysis	19
1.1 Motivation and Research Objectives	20
1.2 Scientific Contributions and Thesis Structure	22
1.3 Publications	26
I Setting the Stage	33
2 Communication Analysis: A Conceptual Framework	35
2.1 Formalizing Communication Analysis	36
2.2 A Short History of Communication Analysis	38
2.3 Framework Methodology	40
2.4 Conceptual Framework	42
2.4.1 Input: Data and Information	44
2.4.2 Processing and Models	47
2.4.3 Visual Interface	49
2.4.4 Knowledge Generation	52
2.5 Discussion and Future Work	54
2.5.1 Expert Validation	54
2.5.2 Survey Findings and Research Opportunities	55
2.5.3 Limitations and Future Work	59
2.6 Conclusion	60
3 Ethical Considerations: Potentials, Limits, and Implications	63
3.1 The Ethical Dimension	65
3.2 Ethical Challenges for Communication Analysis	67
3.3 Scenario Analysis	69
3.3.1 Digital Analysis and Employed Technology	69
3.3.2 The PEGASUS Research Project	70
3.3.3 A Scenario in Police Intelligence Work	71
3.3.4 Conflicts of Interest between Different Stakeholders	72

3.4	Mitigation Techniques through Visual Analytics	75
3.4.1	Technical Measures	75
3.4.2	Balancing Advantages and Risks Through the Interface	80
3.4.3	Risks, Limitations and Additional Measures Required	81
3.4.4	Negotiating Risks and Advantages Through the Interface	83
3.5	Conclusion	84
II	Identification and Interpretation	87
4	Communication Pattern Identification and Interpretation: A Survey of Hypergraph Model Visualizations	89
4.1	Hypergraph Visualizations	90
4.2	Related Work	91
4.3	Methodology	92
4.3.1	Source and Selection Methodology	92
4.3.2	Comparison Framework	93
4.4	Literature Survey	94
4.4.1	Node-Link-based Approaches	94
4.4.2	Timeline-based Approach	97
4.4.3	Matrix-based Approach	97
4.5	Comparison and Analysis	98
4.6	Discussion and Future Work	100
4.7	Conclusion	101
5	Communication Pattern Identification using Hyper-Matrix	103
5.1	Introduction to Temporal Hypergraphs	104
5.2	Related Work	107
5.2.1	Visualization of Hypergraphs	108
5.2.2	Machine Learning for Hypergraph Models	110
5.2.3	Tasks for Evaluation of Temporal Hypergraph Models	110
5.3	Extension of Machine Learning to Hypergraphs	111
5.3.1	Notation and Formulation of a Temporal Hypergraph	111
5.3.2	Relevance Feedback to the Deep Learning Model	113
5.4	Interactive Hypergraph Model Exploration	114
5.4.1	Model Visualization	114
5.4.2	Interactive Exploration and Drill-Down	116
5.4.3	Visual Analytics for Model Updates	119
5.5	Case Study: Forum Communication Data	120
5.6	Formative Evaluation	123
5.6.1	Study Procedure	123
5.6.2	Findings and Lessons Learned	124
5.7	Discussion and Future Work	127
5.8	Conclusion	128
6	Communication Pattern Interpretation through Conversational Dynamics	131
6.1	Communication as Meta-Data Events	132

6.2	Related Work	133
6.3	Communication Behavior Modeling	134
6.4	Visual Analytics of Conversational Dynamics	137
6.5	Conclusion	140

III Holistic Approaches 143

7 CommAID: A Text-based Visual Analytics Technique 145

7.1	The Need for Holistic Communication Analysis	146
7.2	Related Work	148
7.3	Challenges and Design Decisions	151
7.4	Methodology: Modeling Communication	152
7.4.1	Abstract Level	152
7.4.2	Standard Task Levels	153
7.4.3	Dynamics Level	153
7.4.4	Thematic Level	154
7.5	System Design	154
7.5.1	CommAID Interface Design	156
7.5.2	Level Properties	159
7.5.3	Machine-Learning-based Retrieval System	160
7.5.4	Analytical Provenance	161
7.6	Evaluation	162
7.6.1	Formative Domain Expert Assessment	162
7.6.2	Case Study	165
7.6.3	User Evaluation	165
7.7	Discussion and Future Work	166
7.8	Conclusion	167

8 MULTI-CASE: A Multimodal Visual Analytics Technique 171

8.1	Multimodal Communication Analysis	172
8.2	Related Work	176
8.3	Methodology: Model Development	178
8.4	System Design	181
8.4.1	Data Model	182
8.4.2	Component Integration	183
8.4.3	Interfaces and Interaction Principles	183
8.5	Evaluation	186
8.5.1	Case Study	187
8.5.2	Ethics Design Guidelines	189
8.5.3	Intelligence Capability Assessment	190
8.5.4	Model Evaluation	193
8.5.5	Domain Expert Evaluation	193
8.6	Discussion and Future Work	198
8.6.1	Findings and Lessons Learned	198
8.6.2	Limitations and Future Work	200
8.7	Conclusion	201

IV Closing	205
9 Conclusion	207
9.1 Summary	207
9.2 Future Research Aspects	209
9.3 Concluding Remarks	211
Bibliography	213

List of Figures

1.1	Overview of the dissertation structure	24
2.1	Paper collection and coding process steps	41
2.2	Main dimensions of the design space	43
2.3	Conceptual framework of communication analysis systems	45
3.1	The stakeholders involved in communication analysis from the perspective of intelligence analysis	73
3.2	Example for the interactive exploration of a deep learning model	77
3.3	Examples of active learning in Conversational Dynamics and Hyper-Matrix	79
3.4	Example sketch of a provenance history component using a directed acyclic graph approach	80
4.1	The paper collection process	93
4.2	Node-link-based approaches for (temporal) hypergraph visualizations.	95
4.3	Timeline-based approaches for (temporal) hypergraph visualizations.	97
4.4	Matrix-based approaches for (temporal) hypergraph visualizations.	98
5.1	Hyper-Matrix, a novel approach to explore and refine temporal hypergraph models using visual analytics using interactive multi-level matrix-based visualizations	106
5.2	High-level workflow of Hyper-Matrix, showcasing the main components and the interaction flow for the exploration and refinement of temporal hypergraph models	112
5.3	Semantic zoom levels and the different filtering levels leveraged	115
5.4	Comparison of different matrix reordering techniques to facilitate the detection of similar groups and connections	118
5.5	Resulting changes in the model prediction after requesting a model update	120
5.6	Case study comparison of different approaches	122
6.1	Visual analytics workflow of conversational dynamics	133
6.2	Individual communication events are represented as a communication density using kernel density estimation	135
6.3	Communication Sequence visualization showing filtered communication episodes on a vertical timeline	137

6.4	By providing feedback for some data samples, users train ML models to identify relevant conversational dynamics in episodes	138
7.1	High-level workflow of Hyper-Matrix, highlighting the main components and the interaction flow for the communication network analysis	155
7.2	CommAID, an integrated visual analytics technique to analyze communication networks through dynamics modeling, semantic pattern retrieval, and an interactive multi-level matrix-based visualization . .	157
7.3	Overview of the three views in CommAID.	157
7.4	Overview of the details-on-demand offered by different semantic zoom levels	158
7.5	Illustration for the step-by-step process for creating a named entity relation pattern query	159
7.6	By providing feedback, users train ML models to identify relevant conversational dynamics in episodes	161
7.7	Provenance history component showing previous filter states as nodes in a directed acyclic graph	161
8.1	MULTI-CASE: A holistic visual analytics framework tailored towards ethics-aware and multimodal intelligence exploration	175
8.2	High-level architecture of MULTI-CASE	180
8.3	Neighborhood exploration	184

List of Tables

2.1	Paper venues and analysis statistics	41
2.2	Communication analysis system classification	50
2.2	Legend for Communication analysis system classification (cont.) . . .	51
4.1	Comparison of hypergraph visualization techniques	99
7.1	The available analysis levels of the CommAID system	156
8.1	Accuracy on the validation dataset for the baselines and our model .	194

The principal impulse by which I was directed was the earnest endeavor to comprehend the phenomena of physical objects in their general connection, and to represent nature as one great whole, moved and animated by internal forces.

— Alexander von Humboldt, *Explorer*



Introduction to Communication Analysis

Communication is something inherently natural to all of us—and we all have some kind of ingrained conceptualization and understanding of what communication does mean to us personally. We all do communicate constantly: We write, we speak, and we use gestures and facial expressions virtually every day, to name just a few examples of communication behavior. Communication, of course, also happens with and between animals and, in fact, between many different kinds of living organisms in a variety of ways. It forms the basis for group behavior; it is essential to facilitate the buildup of civilization and underpins the inner workings of communities. As such, the definition of the term **communication** is not as clear cut, and its precise definition slightly varies over time and between disciplines: The etymology of the term is rooted in the Latin verb *communicare*, which can be roughly translated as *sharing, transferring, or making common*. As such, it describes the complex process of participating in an exchange of information between individuals, often embedded in a particular context as well as the interpretation of said information to some form or meaning [HW13].

Communication as a **broader topic** has been studied for millennia, often with a focus on rhetorics, with many famous works produced in Ancient Greece and Rome by scholars such as Socrates, Plato (e.g., *Gorgias* [Pla11]), or Cicero (e.g., *De Oratore* [Cic76]). More modern studies on *communication as a process* began in the 20th century, where researchers such as Charles Cooley [Coo09] are credited with starting

Contents

1.1	Motivation and Research Objectives	20
1.2	Scientific Contributions and Thesis Structure	22
1.3	Publications	26

a more systematic research direction, and the origins of human communication have been explored extensively [Tom08]. Nowadays, communication behavior is studied in many different disciplines, such as social psychology, behavioral biology, linguistics, and the digital humanities, with an accompanying large body of research. However, this research has not been compiled and considered in a coherent framework in the context of digital communication and its semi-automated analysis in computer science, which is the focus of this dissertation.

1.1 Motivation and Research Objectives

Over the past two decades, human interactions have seen a significant shift towards digital media and modalities [NGF⁺20]. This increasing prevalence of *digital communication* has led to a growing interest as well as need in the exploration of human communication data across a diverse set of research areas, from digital humanities to business sectors to national security applications. Advances in computer science, and most notably in the field of **big data analytics**, have facilitated this endeavor by offering novel opportunities to explore and extract meaningful insights and knowledge from vast amounts of data. As a result, there has been a proliferation of sophisticated digital systems designed to analyze such data. This development was primarily prompted by the complexities involved in manually (or supported only by basic analytical tools) managing such extensive, heterogeneous, and multifaceted data sets. Innovative methods have been developed to address these challenges more effectively: These entail complex big data analysis techniques like meta-data analysis [MHVB13], pattern recognition [HDL⁺09], social network analysis [Sco17], and natural language processing [EGJ⁺16]. All these approaches are often assisted through machine learning [FAS⁺20].

Despite the considerable volume of **prior research** focused on these *individual* aspects of (semi-) automated communication analysis, the integrated and interactive analysis has received comparatively less attention. However, the increasingly complex nature of communication data poses significant challenges to the process of knowledge extraction. This necessitates further research into this cross-disciplinary area. Challenging is that the field remains unformalized, and the disparate results stemming from the individual analysis modalities can yield contradictory findings, making their comparison difficult. This is because relevance tends to be highly subjective and heavily reliant on specific users, domains, and tasks while experts grapple with articulating and formalizing their constraints into algorithmic terminology a priori. Moreover, without automation, integrating and evaluating the results within a shared context remains an almost insurmountable task. Each approach is limited to its local context, hindering support for cross-matching and inter-modal analysis. Considering the case of investigative journalists, analyzing

transcripts obtained from confidential sources, applying different methodologies in isolation might yield conflicting or inconsistent information. Yet, when considered in unison, these various methodologies could paint a more comprehensive and meaningful picture, revealing valuable insights. Additionally, the previous expertise expressed through domain knowledge and contextual interpretation cannot be understated, playing a significant role. *Classical communication research*, primarily in psychology, has demonstrated that the evaluation of diverse communication modalities [McL64; WBJ74; Sch81] can lead to disparate results, and, therefore, a holistic, context-oriented approach to communication analysis is recommended as opposed to isolated analysis. As such, analysts working with communication data often cannot solely rely on completely automated, single-paradigm solutions. Instead, they require techniques that support a broader analysis and provide them—as human operators—with some agency that (in combination) can considerably improve interpretation.

Visual analytics [KAF⁺08; KKEM10] combines machine learning methodologies and interactive visual interfaces to enable and enhance human sense- and decision-making. The field is characterized by the combination of computational data analysis and multi-faceted, interactive data visualizations, which are tightly coupled through rapid feedback loops, leveraging interactive human sense-making and intuition. The aim is to seamlessly integrate human cognitive and decision-making processes with computational capabilities, facilitated through an iterative, frequent feedback loop. This iterative analysis process forms the core of visual analytics and can be expressed as an explicit process model for knowledge generation [SSS⁺14]. Based on research in human sense-making, the information visualization pipeline [CMS99], and the knowledge discovery process in databases [FPM92], this model illustrates how a user receives support at each step, starting from exploratory analysis, proceeding to hypothesis verification (confirmatory), and finally arriving at knowledge generation. Thereby, visual analytics is better suited for handling ill-defined or open-ended tasks—which are frequently encountered in communication analysis across fields such as intelligence or investigative journalism—while surpassing the capabilities of fully automated systems. With this human supervision and oversight, visual analytics proves highly effective for the design and operation of interactive communication analysis systems that consider multiple information sources and modalities in context, taking into account ethical and privacy aspects, simplifying implementation details, and facilitating an interactive and iterative knowledge generation process, while maintaining high levels of explainability and transparency.

This **dissertation** considers how modern technology facilitates a comprehensive and holistic analysis of human communication. A particular focus is given to investigative domains, like intelligence and investigative journalism, but very similar use cases equally occur in business environments. We consider previous

research on communication analysis from other disciplines but adapt, transfer, and apply the results to discuss relevant aspects of digital interaction from a computer science perspective. The result is a coherent framework in the context of digital communication and its semi-automated analysis that encompasses both aspects of social psychology and technical considerations. As part of this dissertation, we investigate existing approaches, consider and discuss important ethical and privacy considerations, and develop several techniques and approaches to improve communication analysis through the use of interactive visualization methods for steering the analysis while incorporating domain knowledge. Thus, this dissertation focuses on these challenges and its main focus can be summarized in the following three **research questions**:

- **(RQ1)** *What is the potential role of visual analytics techniques in enhancing the interactive analysis and understanding of human communication data?*
- **(RQ2)** *What ethical challenges are faced by communication analysis in a digital age, and what role can visual analytics play in addressing them?*
- **(RQ3)** *What are the challenges and potential strategies for using visual analytics techniques to investigate human communication data, foster interdisciplinary analysis, and support investigations across domain boundaries?*

1.2 Scientific Contributions and Thesis Structure

This dissertation contributes, develops, discusses, and evaluates visual analytics methods for analyzing digital human communication, considering it a holistic and multimodal challenge. To achieve this goal, the following contributions are made:

- A **state-of-the-art survey** and comparison of existing and related approaches of communication analysis systems. From this survey, as well as traditional communication research and technical considerations, we derive a **conceptual framework** of communication analysis. This overview allows us to identify open challenges and **research opportunities**, enabling us to discuss **RQ1**, while addressing them within this dissertation ([Chapter 2](#))
- A detailed analysis on the **ethical frictions and tensions** involved in communication analysis, in particular in the intelligence domain, as well as a **scenario-based stakeholder analysis** of the different actors involved. Based on this, we perform a **critical reflection** on how visual analytics solutions can foster ethical and privacy awareness, negotiate between the inherent trade-offs as an integral part of the design process, and support communication analysis in investigative domains, thereby further investigating **RQ1** and primarily addressing **RQ2** ([Chapter 3](#)).

- To address the **identification of communication** and communication participants, and thereby the *first* part of **RQ3**, we initially survey hypergraph-based model visualization techniques before we present a novel technique, HYPER-MATRIX, to explore and refine underlying temporal hypergraph prediction models through the use of an interactive framework. It encompasses a new, multi-level matrix-based approach and a tight coupling through a relevance feedback loop to integrate and leverage user knowledge for the exploration process. The approach is evaluated through a case study, and the prototype is assessed through a formative expert evaluation (Chapter 4 and Chapter 5).
- To **interpret communication** behavior, we present and evaluate a technique, Conversational Dynamics, presenting a different potential strategy to investigate human communication as raised with the *first* part of **RQ3**, proposing to model communication based on its metadata, detect individual communication episodes, and define features to characterize the resulting communication behavior between entities. The feasibility of this approach is presented through a prototype and an expert interview (Chapter 6).
- To facilitate a more **holistic approach** to communication analysis, addressing the *second* part of **RQ3**, we present a blueprint for a novel, interactive framework, COMMAID, that enables a combined network and content analysis, primarily focusing on text data. We discuss the challenges and design choices involved, describe a case study, as well as evaluate the prototype through an expert user evaluation (Chapter 7).
- Finally, we contribute MULTI-CASE, a **holistic and multimodal visual exploration framework** for intelligence analysis, further rounding off the answer to the *second* part of **RQ3**. The framework is based on the lessons learned in the previous chapters. It covers some of the shortcomings identified in current research, encompasses ethical and privacy considerations, and facilitates a joint agency between the underlying machine learning (transformer) model and the human user. We conduct an extensive evaluation of this technique based on all the aspects considered before, describe a case study scenario, and evaluate the prototype as part of an expert evaluation (Chapter 8).

For the purpose of this dissertation, the contributions are logically structured to build upon each other, starting with a formalization of the field and general considerations before investigating individual approaches and then combining the lessons learned in overarching techniques. The overall **conceptual structure** is described in Figure 1.1, which aligns with the individual chapters:

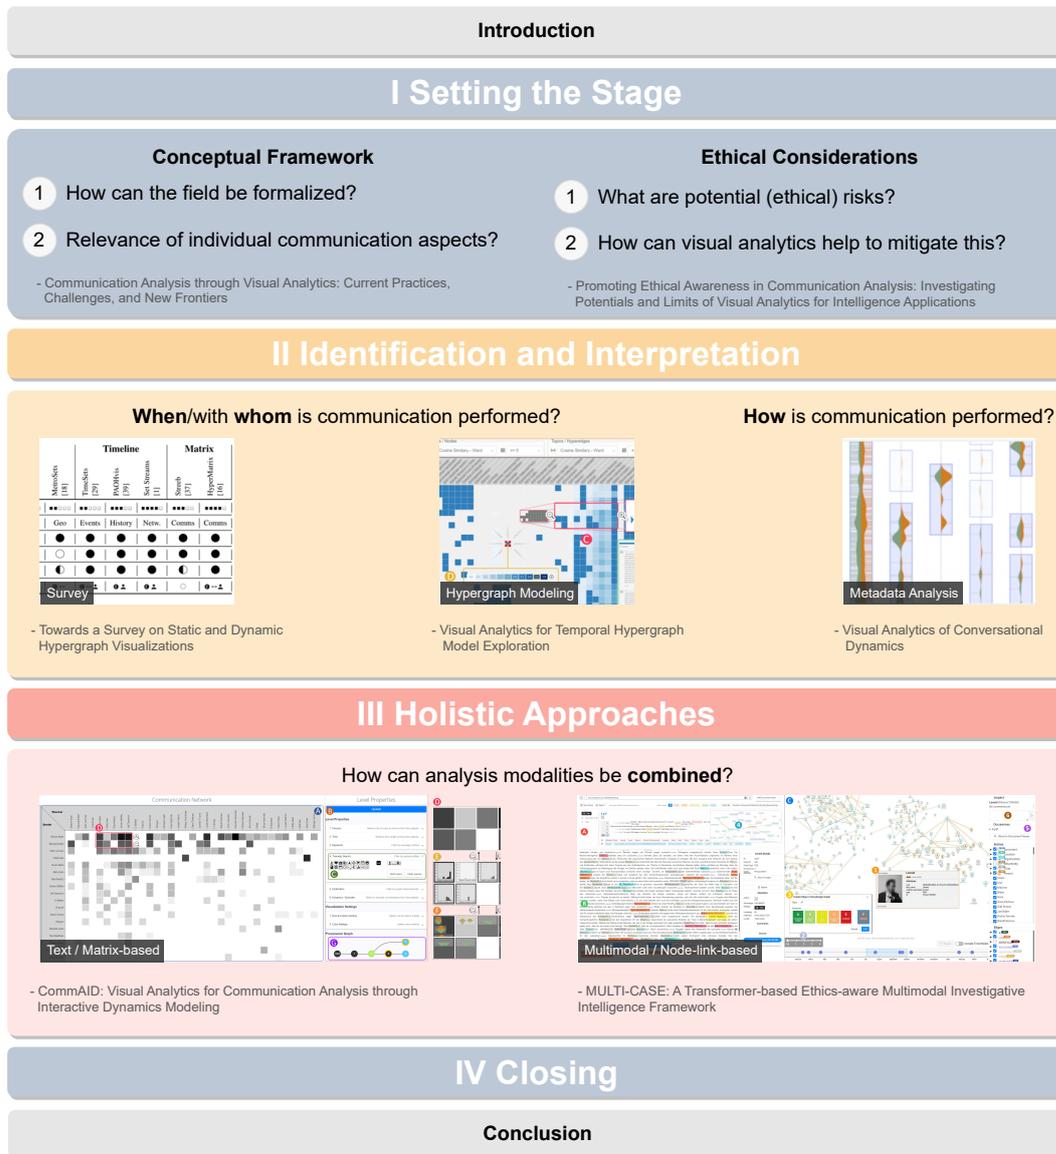


Figure 1.1: Overview of the conceptual structure of this dissertation: This dissertation focuses on how human communication analysis can be performed through semi-automated analysis in the context of computer science. The work is divided into four parts: After the introduction, the Part I sets the stage, covering the conceptual framework, evaluating the state-of-the-art, formalizing the field as well as discussing ethical considerations while proposing how visual analytics can support communication analysis. In Part II, different techniques are discussed that relate to the identification and interpretation of communication. We survey hypergraph model visualizations and present a technique for the identification of communication, as well as metadata pattern analysis for its interpretation. In Part III, we discuss holistic approaches, first on a primarily text-based level and then from a multimodal perspective to communication analysis. The conclusions in the closing Part IV summarizes the results and gives an outlook on the field.

After this introduction, Chapter 2 positions this dissertation within the context of related work and the state-of-the-art while outlining the research gaps it fills. Chapter 3 investigates ethical and privacy risks of communication analysis and shows ways in which visual analytics can address them. Chapter 4 and Chapter 5 consider the identification of communication patterns, while Chapter 6 focuses on their interpretation. The following two Chapter 7 and Chapter 8 both present a holistic and a multimodal framework that combines and summarizes the lessons learned in this dissertation. Finally, Chapter 9 summarizes and concludes this dissertation and presents an overview of further research perspectives.

Citation Rules and Good Scientific Practices

At the beginning of each chapter in this dissertation, I explicitly state the publications this chapter is based upon and the origins of the texts. To avoid any plagiarism and unintended self-plagiarism, I use the following notation to refer to which content I use from the corresponding papers and in which form.

- Sections that are *"taken from"* the listed publications have been largely copied from the corresponding papers and contain only minor structural or literal changes, if not indicated otherwise. In these publications, I was responsible for the main contributions as well as writing virtually the entire paper.
- Sections that are *"based on"* the listed publications are mostly restructured versions of the corresponding papers, and some content has been modified. In these publications, I was responsible for the main contributions as well as writing virtually the entire paper. However, the contributions have been adapted, restructured, or extended to fit nicely into this dissertation.
- A few aspects in this dissertation rely on the groundwork of—or are easier to understand in context with—contributions of other authors that were originally part of joint publications, and where I cannot claim *sole* authorship of idea and text. When absolutely necessary for the storyline in this work, I have—with their permission—paraphrased and summarized their findings in this dissertation, based on the publication record. This only concerns a few minor subsections, namely subsections 3.2, 3.4.2, 3.4.3, and 3.4.4 based on co-author contributions in [FHJ⁺22] as well as subsections 5.2.2 and 5.3.2 based on co-author contributions in [FAS⁺20]. For details of the contributions, in particular regarding these subsections, check the contribution clarification in Section 1.3.

All ideas and concepts presented in this thesis originate from me, if not otherwise indicated, or are quoted appropriately. To be as transparent as possible regarding the emerging use of digital tools and good scientific practices, their use did *not* play a role for the ideas and content of this thesis and was strictly limited to

stylistic, phrasing, and corrective actions (e.g., spell-checking, style- and phrasing suggestions, rewording, translation), operating on *existing*, *self-written* texts and all tool suggestions were manually verified and adapted.

1.3 Publications

During my time as a doctoral student and research associate, I have contributed to and been part of several scientific works, which have been published in international peer-reviewed journals and conference proceedings, some of which serve as the basis for this dissertation.

Thesis-Relevant Publications and Contribution Clarification

The following list gives an overview of the publications that form the basis of this thesis and the division of work between the authors. The publications are sorted by their appearance in the following chapters.

- [FDS⁺22b]: **Maximilian T. Fischer**, Frederik Dennig, Daniel Seebacher, Daniel A. Keim, and Mennatallah El-Assady. “Communication Analysis through Visual Analytics: Current Practices, Challenges, and New Frontiers”. In: *2022 IEEE Visualization in Data Science (VDS)*. 2022, pp. 6–16. ISBN: 978-1-6654-5721-7. DOI: **10.1109/VDS57266.2022.00006**

Contribution clarification: This paper originates from a discussion between Menna and myself. I took the project lead, coming up with the overall structure. Further, I identified the basic research question as well as contributions and provided background on related work. The initial survey collection and the construction of the conceptual framework, its findings, and their discussion were also done by me. Frederik Dennig and Daniel Seebacher helped me in that task, also in particular in giving feedback and suggestions for multiple iterations of the framework. Daniel Keim and Menna El-Assady provided comments on the paper drafts, and all helped in revising the sections through helpful comments. I wrote all sections myself and revised suggestions from my co-authors several times. Thus, I will reuse the paper text in [Chapter 2](#).

- [FHJ⁺22]: **Maximilian T. Fischer**, Simon D. Hirsbrunner, Wolfgang Jentner, Matthias Miller, Daniel A. Keim, and Paula Helm. “Promoting Ethical Awareness in Communication Analysis: Investigating Potentials and Limits of Visual Analytics for Intelligence Applications”. In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’22)*. Association for Computing Machinery, 2022, pp. 877–889. ISBN: 978-1-4503-9352-2. DOI: **10.1145/3531146.3533151**

Contribution clarification: This paper was a close collaboration between

Simon David Hirsbrunner and Paula Helm from the IZEW (Internationales Zentrum für Ethik in den Wissenschaften, i.e., International Center for Ethics in the Sciences and Humanities) located at the University of Tübingen. The original idea for this paper was based on a discussion I had with several of my students during a week-long academy (Roggenburg A) which I was offered to teach as part of the Studienstiftung des Deutschen Volkes academy cycle in the spring of 2021. I initially took the lead on this paper, coming up with the overall structure and content topics, as well as being responsible for the background, scenario, stakeholder, and technical considerations. In particular, my contributions focus on the introduction (Section 1), the scenario analysis (Section 3), together with the stakeholder analysis, the technical measures (Section 4.1), and the conclusion (Section 5), which I all wrote as initial drafts. Simon David Hirsbrunner, Paula Helm, and I worked together on the ethical dimensions (Section 2) as well as the advantages, risks, and interfacing debates (Sections 4.2-4.4) as initial drafts. I reworked those sections again to streamline, extend, and sharpen them, adding technical considerations and backlinks, and both also provided helpful suggestions for my draft parts. Wolfgang Jentner, Matthias Miller, and Daniel Keim provided comments on the paper drafts and some conceptual discussions. Accordingly, I will reuse the paper text for Sections 1, 3, 4.1, and 5 while summarizing Sections 2 and 4.2-4.4 as Sections 3.2 and 3.4.2, 3.4.3, 3.4.4 as part of [Chapter 3](#).

- [FFKS21]: **Maximilian T. Fischer**, Alexander Frings, Daniel A. Keim, and Daniel Seebacher. “Towards a Survey on Static and Dynamic Hypergraph Visualizations”. In: *2021 IEEE Visualization Conference (VIS)*. IEEE, 2021, pp. 81–85. DOI: [10.1109/VIS49827.2021.9623305](https://doi.org/10.1109/VIS49827.2021.9623305)

Contribution clarification: This paper was a close collaboration between Alexander Frings and myself as part of a seminar (I proposed, sketched, and supervised his seminar work). I defined the initial research objective and contributions while providing a starting ground. Based on the initial literature search conducted by Alexander Frings under my guidance and feedback, he summarized his search results in the seminar before we designed the final evaluation criteria used in this paper. I later extended these criteria and also added some additional approaches that were not uncovered during the initial literature search. Daniel Seebacher and Daniel Keim provided feedback on the general idea and commented on paper drafts. I proposed and closely oversaw the research process, wrote the major parts of the text based on the search conducted by Alexander Frings, and extended them with my own ideas as discussed above, while I also revised and extended suggestions by Alexander Frings several times during the writing process. Thus, I reuse the major content in [Chapter 4](#).

- [FAS*20]: **Maximilian T. Fischer**, Devanshu Arya, Dirk Streeb, Daniel Seebacher, Daniel A. Keim, and Marcel Worrying. “Visual Analytics for Temporal Hypergraph Model Exploration”. In: *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2020), pp. 550–560. DOI: [10.1109/TVCG.2020.3030408](https://doi.org/10.1109/TVCG.2020.3030408)

Contribution clarification: This paper originated from a close collaboration and joint work with Devanshu Arya and Marcel Worrying from the University of Amsterdam (UVA) and was based on an initial discussion with Dirk Streeb. I took the project lead and came up with the framework idea, and designed the visualization and interactive prototype. The whole interface and backend system was developed by me, except that Devanshu Arya contributed the re-trainable geometric machine learning model as a plugin. I also designed the case study and performed the formative evaluation. I wrote the large majority of the sections myself, namely the introduction (Section 1), the Related Work (Section 2, except 2.2), the first Section on Machine Learning for Hypergraphs (Section 3.1), the system design (Section 4), the case study (Section 5), the formative evaluation (Section 6) with the findings, the discussion (Section 7), as well as the conclusion (Section 8). Devanshu Arya contributed a background on geometric deep learning (Paper Section 2.2) as well as the extension of previous work for interactive retraining (Paper Section 3.2). Dirk Streeb, Daniel Seebacher, and Marcel Worrying provided comments on the paper drafts and helped in revising and streamlining the sections, while Daniel Keim provided feedback on the approach and advised on the paper drafts. I wrote all sections myself, with the above-mentioned exceptions. Thus, I will reuse the paper text for all sections, except for the original Sections 2.2 and 3.2, which have been summarized in Sections 5.2.2 and 5.3.2 as part of [Chapter 5](#).

- [SFS*19]: Daniel Seebacher, **Maximilian T. Fischer**, Rita Sevastjanova, Daniel A. Keim, and Mennatallah El-Assady. “Visual Analytics of Conversational Dynamics”. In: *EuroVis Workshop on Visual Analytics (EuroVA)*. ed. by Tatiana von Landesberger and Cagatay Turkyay. EuroVA. Porto, Portugal: The Eurographics Association, 2019. ISBN: 978-3-03868-087-1. DOI: [10.2312/eurova.20191130](https://doi.org/10.2312/eurova.20191130)

Contribution clarification: This work was a continuation, formalization, and evaluation of the overall research direction I started in my Bachelor’s Thesis (*Visual Analytics for Detecting Patterns in Large Communication Networks* [Fis18]) and piqued my interest in my doctorate’s research field. It was a close collaboration with my then advisor Daniel Seebacher, as well as Rita Sevastjanova, Daniel Keim, and Menna El-Assady. I designed and implemented the research prototype and further conducted the expert study and the literature review. Daniel Seebacher provided guidance and supervision for my first paper as well as helpful feedback. Rita Sevastjanova, Daniel Keim, and Menna El-Assady

provided feedback on the general idea and commented on paper drafts. I wrote the major parts of the text and revised all sections several times. Thus, I reuse the text in [Chapter 6](#).

- [FSS⁺21]: **Maximilian T. Fischer**, Daniel Seebacher, Rita Sevastjanova, Daniel A. Keim, and Mennatallah El-Assady. “CommAID: Visual Analytics for Communication Analysis through Interactive Dynamics Modeling”. In: *Computer Graphics Forum* 40.3 (2021), pp. 25–36. ISSN: 01677055. DOI: [10.1111/cgf.14286](https://doi.org/10.1111/cgf.14286)

Contribution clarification: This work was a close collaboration between Daniel Seebacher, Rita Sevastjanova, Daniel Keim, and Menna El-Assady. I took the project lead and came up with the initial idea. Further, I set out the research goals and the contributions, performed the literature review, conducted the domain expert evaluations, and defined the case study. Daniel Seebacher implemented a first draft prototype, which I extended and advanced further. Rita Sevastjanova contributed the named-entity pattern search method, while Daniel Keim and Menna El-Assady provided feedback on the general idea and commented on paper drafts. I wrote all sections myself and revised suggestions from my co-authors several times. Thus, I reuse the text in [Chapter 7](#).

- [FMJ⁺24]: **Maximilian T. Fischer**, Yannick Metz, Lucas Joos, Matthias Miller, and Daniel A. Keim. “MULTI-CASE: A Transformer-based Ethics-aware Multimodal Investigative Intelligence Framework”. In: (2024), pp. 1–16. DOI: [10.48550/arXiv.2401.01955](https://doi.org/10.48550/arXiv.2401.01955)

Contribution clarification: This work is both an extension of the holistic idea started with CommAID, as well as a bringing together of the ideas and approaches presented throughout the thesis. It was a close collaboration between Yannick Metz, Lucas Joos, Matthias Miller, and Daniel Keim. I was responsible for the overall paper idea and research objectives. I conducted the literature review, designed the system’s architecture, wrote the technique description, and performed the comparative evaluations, as well as wrote the case study. Yannick Metz performed the machine learning training. In particular, he wrote and supervised the training scripts. Lucas Joos, Matthias Miller, and Daniel Keim provided feedback on the general idea and commented on the paper drafts. I wrote all sections of the text myself and revised suggestions from my co-authors several times. Thus, I reuse the text in [Chapter 8](#).

Additional Publications

In addition to this, there were a number of publications that I authored or contributed to, which are listed here for additional context, but have not been included in this thesis:

- [FKS19]: **Maximilian T. Fischer**, Daniel A. Keim, and Manuel Stein. “Video-Based Analysis of Soccer Matches”. In: *Proceedings of the 2nd International*

Workshop on Multimedia Content Analysis in Sports. MMSports. New York, NY, USA: ACM, 2019, pp. 1–9. DOI: [10.1145/3347318.3355515](https://doi.org/10.1145/3347318.3355515)

- [DFB⁺21]: Frederik L. Dennig, **Maximilian T. Fischer**, Michael Blumenschein, Johannes Fuchs, Daniel A. Keim, and Evanthia Dimara. “ParSetgnostics: Quality Metrics for Parallel Sets”. In: *Computer Graphics Forum* 40.3 (2021), pp. 375–386. ISSN: 01677055. DOI: [10.1111/cgf.14314](https://doi.org/10.1111/cgf.14314)
- [JSS⁺22]: Wolfgang Jentner, Fabian Sperrle, Daniel Seebacher, Matthias Kraus, Rita Sevastjanova, **Maximilian T. Fischer**, Udo Schlegel, Dirk Streeb, Matthias Miller, Thilo Spinner, et al. “Visualisierung der COVID-19-Inzidenzen und Behandlungskapazitäten mit CoronaVis”. In: *Resilienz und Pandemie*. Ed. by Andreas Hermann Karsten and Stefan Voßschmidt. Stuttgart: Kohlhammer, 2022, pp. 176–185. ISBN: 978-3-17-039930-3
- [NLH⁺22]: Adela Cañete Nieto, Ruth Ladenstein, Barbara Hero, Sabine Taschner-Mandl, Ulrike Pötschger, Vanessa Düster, Martinez De Las Heras, Blanca, Ana Jiménez Pastor, Eva Bozsaky, **Maximilian T. Fischer**, et al. “PRIMAGE - An Artificial Intelligence-based Clinical Decision Support System for Optimized Cancer Diagnosis and Risk Assessment - A Progress Update”. In: *SIOPEN AGM*. Athens, 2022
- [FAMK22]: **Maximilian T. Fischer**, Dennis Ackermann, Yannick Metz, and Daniel A. Keim. “Agent-based Visual Analysis of Public Transport Infrastructure”. In: *2022 Smart Country Convention (SCCON)*. Berlin, 2022
- [JKF⁺23]: Lucas Joos, Karsten Klein, **Maximilian T. Fischer**, Frederik L. Dennig, Daniel A. Keim, and Michael Krone. “Exploring Trajectory Data in Augmented Reality: A Comparative Study of Interaction Modalities”. In: *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2023, pp. 790–799. DOI: [10.1109/ISMAR59233.2023.00094](https://doi.org/10.1109/ISMAR59233.2023.00094)
- [JFKF23]: Lucas Joos, **Maximilian T. Fischer**, Daniel A. Keim, and Johannes Fuchs. “Aesthetic-Driven Navigation for Node-Link Diagrams in VR”. in: *Proceedings of the 2023 ACM Symposium on Spatial User Interaction (SUI)*. vol. 25. Association for Computing Machinery, 2023. DOI: [10.1145/3607822.3614537](https://doi.org/10.1145/3607822.3614537)

Part I | Setting the Stage

The purpose of visualization is insight, not pictures.

— Ben Shneiderman, HCI researcher

2

Communication Analysis: A Conceptual Framework

Communication behavior is studied in many different scientific disciplines. In computer science, traditionally, the automated analysis of digital human communication data often focuses on specific aspects, such as content or network structure, in isolation. This can provide limited perspectives while making cross-methodological analyses difficult, for example, in domains such as investigative journalism. Communication research in psychology and the digital humanities instead stresses the importance of a holistic approach to overcome these limiting factors. However, for many of the research aspects considered in social psychology, it is not inherently obvious how they can be transferred and applied in computer science or even automated as part of analysis processes.

In the following, we begin to construct a conceptual framework of digital communication analysis that supports a computer science-based discussion and analysis of the topic. We base this framework on communication research, technical considerations, and distinguishing features we observe in existing applications. In the following, after a historical overview of the field, we conduct an extensive survey on the properties of over forty semi-automated approaches that can be considered, to one degree or another, as communication analysis systems. We then investigate how these approaches cover concepts described in theoretical communication research. From these investigations, we derive our design space and contribute a conceptual framework based on communication research, technical considerations, and the surveyed approaches. The framework describes the systems' properties,

Contents

2.1	Formalizing Communication Analysis	36
2.2	A Short History of Communication Analysis . . .	38
2.3	Framework Methodology	40
2.4	Conceptual Framework	42
2.5	Discussion and Future Work	54
2.6	Conclusion	60

capabilities, and composition through a wide range of criteria. These are organized in the dimensions (1) Data, (2) Processing and Models, (3) Visual Interface, and (4) Knowledge Generation. These criteria enable a formalization of digital communication analysis through visual analytics, which, we argue, is uniquely suited for this task by tackling automation complexity while leveraging domain knowledge. With our framework, we identify shortcomings and research challenges, such as group communication dynamics, trust, and privacy considerations (see Chapter 3), and holistic approaches (see Chapter 7 and Chapter 8), among others. Simultaneously, our framework supports the evaluation of systems and promotes the mutual exchange between researchers through a structured, common language, laying the foundations for future research on communication analysis while forming the conceptual basis of this dissertation.

This chapter is based on the publication [FDS⁺22b] and major parts of the following sections have appeared in:

- [FDS⁺22b]: **Maximilian T. Fischer**, Frederik Dennig, Daniel Seebacher, Daniel A. Keim, and Mennatallah El-Assady. “Communication Analysis through Visual Analytics: Current Practices, Challenges, and New Frontiers”. In: *2022 IEEE Visualization in Data Science (VDS)*. 2022, pp. 6–16. ISBN: 978-1-6654-5721-7. DOI: [10.1109/VDS57266.2022.000006](https://doi.org/10.1109/VDS57266.2022.000006).

For a statement of the scientific contributions, as well as the division of responsibilities and work in this publication, please refer to Chapters 1.2 (p. 22ff) and 1.3 (p. 26ff), respectively.

2.1 Formalizing Communication Analysis

Human communication has been fundamentally transformed, especially in the last two decades, becoming increasingly digital, with cost-effective, location-independent, and instant access changing communication behavior. With this transformation to digital communication [Sco09], new research opportunities have emerged in a wide variety of different domains, ranging from engineering to social sciences to business: For example, it has been studied how to visualize the evolution of dynamic communication networks [Tri08], how discourse analysis for digital communication can be enhanced [Her19], or how team communication performance in business settings

can be evaluated [FM08]. For such analyses, digital analysis methods are often used to aid and support the (semi-)manual, domain-specific research methodologies.

In this dissertation, we focus on the field of interactive human communication analysis and specifically on **automated and interactive communication analysis systems** targeting written human communication (in the following: communication analysis systems), most commonly e-mails, chats, or documents. For this dissertation, we define these systems as semi-automated applications that employ visual components for an interactive analysis. We do not consider approaches focusing primarily on a single methodology like sentiment analysis, but those that aim at a cross-methodological analysis among multiple parties. This analysis becomes increasingly relevant in many investigative domains [BISM14; FHJ⁺22].

The research into communication analysis *systems* often lacks [v14; BISM14] **cross-methodological** aspects: the majority of systems focus on either the content of communication *or* on the network aspect *in isolation* instead of considering the fundamental dynamics holistically. This is in contrast to seminal works on human communication research [McL64; WBJ74], recent textbooks [Pea11; McL17], or current communication research in psychology or digital humanities [FM08; Mes09], where often—even when digitally supported [FG14; WCG⁺16]—a holistic view is taken to consider explicit and implicit connotations in context. In contrast, the individual analysis of content, network, and metadata can—for interrelated tasks—lead to an incomplete or biased view, while isolated approaches often introduce discontinuities, increase manual work and hamper cross-methodological detection.

Existing frameworks on digital communication analysis *systems* do not adequately cover this issue due to four reasons: First, the need for such a revised formalization has been recognized [vP18] in communication sciences. So far, the opportunities, challenges, and pitfalls have primarily been described from an application domain-oriented perspective [FG14; WCG⁺16] in the social sciences, while a systematic description is missing, only available for social-media-based approaches [FG14; WCG⁺16]. Ethical considerations [FHJ⁺22] so far play only a small role in the system design. Second, recent efforts have begun to map digital communication systems as a whole [FL20], with a focus on content, infrastructure, and policy aspects, but leaving out the technical considerations, like methods, interfaces, and interaction concepts. Third, the same is true for the classical communication analysis research [McL64; WBJ74; Pea11], which lacks technical considerations and is primed for analog but not digital communication. Fourth, digital communication has also transformed the way we communicate and the modalities we use [OO19], like shorter messages or emoji reactions, requiring an updated framework.

In this work, we want to bridge the **gap** between communication research and modern communication analysis system development. As evident in a few academic works [HDL⁺09; WLY⁺14; KFS⁺19; FSS⁺21; Gro23] and recent commercial systems [Nui20;

Pal20; Dat20], visualization and interactive user steering are a promising way [KAF⁺08; YKSJ07; GKL⁺13; SSS⁺14; Her15; WCG⁺16; CSJ⁺18] to begin to tackle the gap between different analysis modalities. Lack of a common description from both a technical perspective and psychological communication research has made the systematic exploration of the field difficult. This also prevented a broader review of how visual analytics principles are—and could be—employed in communication analysis, how such systems can be categorized, and what a relevant taxonomy would look like. The main objective of this work is to explore these systems from a primarily capability-oriented perspective, in terms of communication research, technical state of the art, and human factors. While we consider and point out these human factors and ethical considerations as much as possible within this framing, we also refer to Chapter 3 for a more detailed background and in-depth discussion on ethical awareness and human factors in communication analysis.

As part of this chapter, we survey state-of-the-art approaches and investigate concepts in communication research to derive a **design space** on communication research, making the following contributions:

Contributions

- The creation of a **conceptual framework** (see Figure 2.3) of communication analysis systems based on communication research, technical considerations, and a systematic review.
- A state-of-the-art **survey** and comparison of existing approaches, assessing their maturity and coverage (see Table 2.2)
- A discussion on the open challenges and implications for future research **opportunities** on communication analysis systems.

With this contribution, we identify research challenges and aid the comparison of approaches while creating a taxonomy for future research on communication analysis through visual analytics.

2.2 A Short History of Communication Analysis

Communication analysis can use a variety of different techniques to analyze communication behavior in its entirety. The origins of **communication research** can be traced back to rhetoric and oratory in ancient times, with the study of rhetoric and oratory as well as persuasion in Ancient Greece and Rome. Many famous works were produced by scholars such as Sokrates, Plato (*Gorgias* [Pla11]), or Cicero (*De Oratore* [Cic76]). The communication process itself started to be studied in the early 20th century, where a more systematic research direction is often credited to researchers such as Charles Cooley [Coo09] (investigating the human nature of

communication), Georg Simmel [Sim08] (studying interaction and group formation in sociology), Walter Lippmann [Lip22] (considering the disparity of expression), and Jacob Moreno [Mor34] (researching human networks and interpersonal connections). The works were extended to communication patterns inside groups (Bevelas [Bav50] and Leavitt [Lea51]) as well as to computer-aided modeling (Shannon [SW49] and Savage and Deutsch [SD60]). From the 1960s onward, the seminal works on media (McLuhan [McL64]), communication theory (Watzlawick et al. [WBJ74]), inter-human communication (von Thun [Sch81]), and communication networks (Roger [RK80]) established the field. It now encompasses diverse techniques [Pea11; McL17], from natural language processing over social network analysis to metadata exploration. All focus on the **human factors** and **communication context**, with channels (Watzlawick et al.) / medium (McLuhan) forming an essential analysis aspect: Be it phrasing or omissions, narrow- or broadcasting to target audiences, subtle implicit messages, expectations of confidentiality (and frankness), or effort in crafting the messages.

With the **advent of digital** processing, the laboursome manual analysis [Ger69] shifted first to digitally supported [LSB02] and later to highly automated analysis. However, it is noticeable [FSS⁺21] that the analysis' completeness developed counter to automation level, with increasing specialization and isolation. For example, modern systems can analyze communication behavior using centrality measures [LZ15] or describe network ties in social sciences [BMBL09]. Specialized visual toolkits have been developed to analyze such networks, like Pajek [BM98] or Gephi [BHJ09]. However, all these approaches primarily focus on the network aspects, omitting most of the meta-data and especially the content. Others focus on content instead, like keyword-based searches [YP04] to filter communication or aim to improve the understanding of communications meaning through sentiment analysis [PL08] or topic modeling [ŘS10].

However, **visual analytics** could support a more comprehensive analysis, as we outline in Part III, discussing the potential for holistic systems. Even existing visual approaches often follow insular approaches, like using node-link-diagrams (e.g., Gephi [BHJ09], and many commercial solutions like IBM's i2 Analyst's Notebook [IBM20], Pajek [BM98], Palantir Gotham [Pal20], DataWalk [Dat20], and Nuix Discover and Nuix Investigate [Nui20]). Another class of approaches uses matrix-based approaches to analyze the communication (or social) relations, for example, MatrixExplorer [HF06] or NodeTrix [HFM07]. Another set of approaches use timeline designs like CloudLines [KBK11], while others like Fu et al. [FHN⁺07] modify graph presentations through multiple planes.

The complexity and ambiguity of the exchanges and modalities [JFDK00] make complete automation difficult, simultaneously raising serious ethical and privacy considerations [FHJ⁺22]. As such, **visual analytics** [KAF⁺08] (for a definition, see Sec-

tion 1.1) is uniquely suited [FSS⁺21; FHJ⁺22] for a holistic approach to communication analysis, considering the subtleties of human communication.

2.3 Framework Methodology

In the following, we aim to tackle the central question of a **common description** of communication analysis:

How can the different approaches in communication analysis systems be described within a common, conceptual framework to allow their mutual comparison?

Framework Basis — We propose to base such a framework on three areas of consideration:

1. The **existing research landscape** of interactive communication analysis systems provides a foundation for the classification of approaches based on measures such as analytical goals, visualization, and interaction methods, or the power of the knowledge generation process.
2. **Communication research** offers decades of research on the particularities of (often non-technical) communication analysis. For this work, we consider concepts from seminal and more recent summary works [SW49; McL64; WBJ74; Sch81; Mes09; Pea11; McL17; CSJ⁺18], described in Section 2.2. Additionally, we study relevant theoretical (non-system) works in computer science, like a survey on text visualization [KK15], group discourse and role analysis [HC15; HC16; FWY⁺18; LPT⁺21], as well as works dealing with semi-manual approaches and user studies, including the human factors (e.g., [JAF10; FNS13; BISM14; SSK⁺16; GF17; MKMS20; Cor19; FHJ⁺22]). However, many of these works miss the transfer from a theoretically analyzed concept to an actual system implementation.
3. **Technical considerations** of the approaches, taking into account design properties such as analyzable data types, data representation, and flow, or limitations like scalability from a technical standpoint.

We discuss the findings from considerations (2) and (3) later in Section 2.4, while (1) requires a broad review:

Existing Research Landscape – Seed Papers — To analyze the state-of-the-art and contribute one angle of classification criteria, we start with a keyword-based seed literature survey. As we aim to inform about the most common ideas in *visual analysis* applications, we restrict the search to high-quality journals and conferences: IEEE Trans. of Vis. and Comp. Graph. (**TVCG** and **IEEE VIS**), Comp.

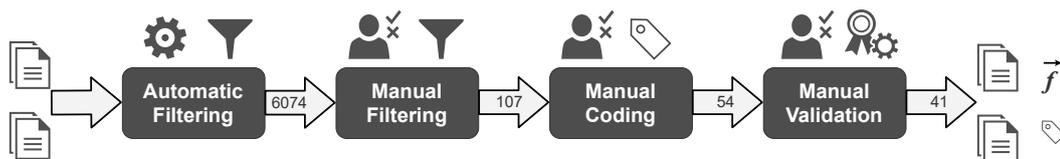


Figure 2.1: Paper collection and coding process steps: (1) Automated filtering, (2) manual filtering, (3) manual coding, (4) manual validation.

Table 2.1: The publications per venue and paper count statistics for each of the steps in the collection and coding process.

Venue	#Collected	#Filtered	#Coded	#Final
IEEE TVCG / IEEE VIS	790	35	27	23
Computer Graphics Forum	495	17	11	8
CHI Proceedings	4789	49	10	4
Commercial Systems	-	6	6	6
Total	6074	107	54	41

Graph. Forum (**EuroVis**), Proc. of the **CHI** Conference on Human Factors in Comp. Sys. and their co-located events. We focus on more recent solutions that can leverage technological advances in the last 15 years (i.e., beginning in mid-2007). We are aware that thereby we exclude some useful techniques (e.g., [FHN⁺07]). Outside the VIS community, in journals like Digital Investigation, only few visualization approaches (e.g., E-Mail Forensics [HDL⁺09]) have been published. Due to the absence of novel visualizations or integrations, they were not included.

Selection Methodology — For the actual paper selection methodology, we follow a four-step approach (see also Figure 2.1). For context, the publications per venue and paper counts analyzed in each step are detailed in Table 2.1. First, we conducted a keyword-based seed search for the words *communication* and *analysis* on the titles, abstract, index terms, and contents of publications in each of the venues described above. Secondly, we went through all these papers’ titles and abstracts manually, discarding those which clearly are not concerned with communication analysis systems, reducing the selection significantly. For CHI, the high number of initial approaches and the high discard rate is due to the abundant use of the phrase *communication* when referring to user actions. In this step, we included approaches suggested by the domain experts. Third, we manually looked at the remaining papers and decided whether they describe a communication analysis system (as defined above) or list it as a potential application. In the final step, we validated the results by checking borderline cases, consequently removing seven papers. Our final collection includes 41 approaches.

Domain Expert Consultation — To broaden the perspective, we consulted with eight domain experts by conducting an informal interview. The experts belong to the

field of law enforcement, working for various European law enforcement agencies, and each has extensive experience with digital investigations, including communication analysis, working in the field from ten to over 30 years. They contributed **(1)** a collection of six **approaches used in practice** [BM98; BHJ09; IBM20; Pal20; Dat20; Nui20] (including commercial) as well as **(2)** insights on **their analytical needs** and **perceived challenges**. The proposed approaches were included in the Selection Methodology from Step 2 onward, and the analytical needs were considered for the classification criteria in Section 2.4. Further, we recruited **additional domain experts** for an interview to **evaluate the completed conceptual framework**, which we discuss in Section 2.5.1.

Regarding challenges, the experts consider it unlikely that an autonomous system can completely replace an experienced-saturated investigator with years of domain-specific knowledge [FHJ⁺22] except in the narrowest or specialized of tasks. As soon as incomplete information is involved and decisions under uncertainty have to be taken, the analysts often follow their hunches, exploring different options, but having difficulty in articulating their reasoning [FSS⁺21]. They explore related and connected information, which they consider important for contextual information [FSS⁺21]. As such, they are used to - and strongly prefer - visually-interactive tools for investigations, as it supports their understanding through rapid-feedback mechanisms [FSS⁺21], increasing their trust [FHJ⁺22]. Such systems have been increasingly deployed in fields such as investigative journalism or criminal investigations [FHJ⁺22]. Nevertheless, many experts are open to new developments and consider systems their companions, supporting them without patronizing or limiting them [FHJ⁺22], relieving them of labor-some manual work. However, they have to be developed with analysts in mind [FHJ⁺22], otherwise potentially overwhelming them or missing key functionality [FAS⁺20]. Black box AI models are received critically except for hints, as the domain experts are often no AI experts, lacking opacity and making it difficult to prove provenance and a chain-of-reasoning that fulfills moral or legal obligations [FHJ⁺22]. Developing systems fulfilling these requirements while leveraging reliable XAI methods are key challenges.

2.4 Conceptual Framework

In the following section, we aim to construct a framework that encompasses discerning aspects of communication analysis systems. As with any taxonomy, the framework is *one* possible version of a taxonomy, developed in several iterations. We justify our considerations overall and for each property, either based on the domain experts' requirements or referencing relevant work when applicable. The overall framework was constructed by the authors collaboratively through the collection of aspects from the three consideration areas (see Section 2.3) and the iterative

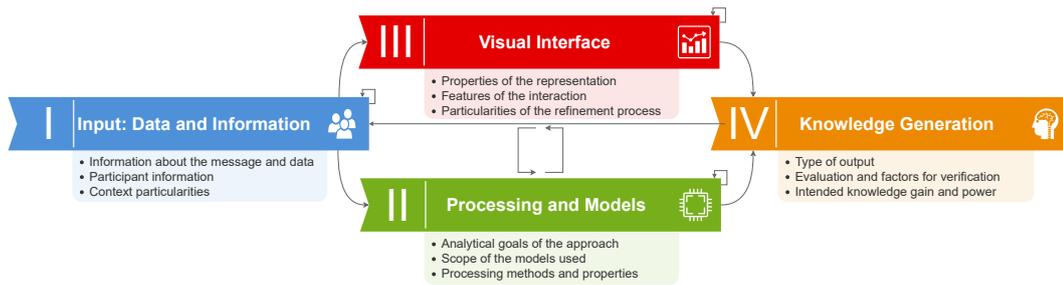


Figure 2.2: Characterization of the four main dimensions of our conceptual framework for communication analysis systems, in the form of a concrete application of the *visual analytics process model* by Keim et al. [KAF⁺08] to the communication analysis domain.

construction of a mutually exclusive and collectively exhaustive description, aiming to group related aspects and align with established nomenclatures. We caution that many presented properties are multifaceted, and our considerations can benefit from community discussion. For the complete conceptual framework, see Figure 2.3 as well as the main dimensions in Figure 2.2, while for the full classification of the surveyed approaches, consult Table 2.2.

Main Considerations – One standard **structuring methodology** is to use a task-based grouping [CSJ⁺18; FAS⁺20]. However, sometimes very different methods are employed for the same task: for example, to discover key persons in a communication network, SNA-based [GKL⁺13] approaches using centrality measures and node-link visualizations are equally applicable as interactively visualized geometric deep learning models [FAS⁺20]. However, both approaches have very different side effects, visualization, and interaction techniques and make very different assumptions about the data. Instead, we follow the second large methodology, a thematic-based grouping, and use a property, representation, and methods-based taxonomy [KK15].

As our **primary goal** is to design a conceptual framework of communication analysis through *visual analytics*, we motivate the main areas by Keim et al.’s established process model [KAF⁺08], but develop each of the four areas specifically for communication analysis using considerations from communication research and our survey. We chose Keim et al.’s model because it is considered the most widely used visual analytics model [SSS⁺14], compared to similar models like Green’s Human Cognition Model [GRF08] or van Wijk’s Visualization Model [van05]. Further, having similar dimensions, the differences between these models are slim from a compartmentalizing perspective. We slightly modify Keim et al.’s terminology, proposing **four main dimensions**, characterized further in Figure 2.2 and the following sections: (I) Input: Data and Information (2.4.1) encompasses the (inferred) content and context with respect to communication research, (II) Processing and Models (2.4.2) discusses the analytical goals and scopes of the systems, (III) Visual Interface

(2.4.3) presents visual and interaction techniques employed, and (IV) Knowledge Generation (2.4.4) discusses the information flow.

2.4.1 Input: Data and Information

This category focuses on information, context, and environment of the communication, in particular theoretical aspects and data properties, with the structuring partly based on classical communication research [SW49; McL64; WBJ74; Sch81; Mes09; Pea11; McL17; CSJ⁺18]. Therefore, we are discussing the content and meaning, context, and relationship aspects of communication extensively. Building on established frameworks [SW49; WBJ74; Sch81], we propose to focus on three interrelated areas: the information as *message*, the *communication participants*, and the *environment* (or context).

Message

The *message* [Sch81] (also central channel [SW49] or content [WBJ74]) refers to the entailed information. From a system's perspective the distinction by *data type* is obvious, while the information can be considered from its actually transported content (*coding* [WBJ74]) and its orthogonal interpretation (*expression* levels [Sch81]):

Data Type — The data type refers to the content type from a technical point of view. When looking at data classification in information visualization [CMS99], we can identify several data types which are relevant for communication analysis: *text*, *audio*, *image*, *video*, and meta-data, related to *network* as well as *time-series*. Based on the usage in current approaches (see Section 2.3), the two most relevant ones are text data (e.g., extracting topics from text [CLT⁺11]) and relation network (structure) data (e.g., social graphs between communication participants [BHJ09]). However, communication can also happen via audio (e.g., telephone or VoIP) or via video chats, comprising audio and moving images, i.e., video data. While our framework focuses primarily on this written (i.e., text) communication, we include these types for completeness. Therefore, it would be possible to include the detection of facial expressions using deep learning [NNVW15] to analyze the analogical code and set it into context (see below). Meta-data in the form of time-series data (e.g., [KBK11], extracting event order and relevance) is often relevant for the communication *context*, regarding regularity and duration.

Coding — The transported meaning is a core aspect of the communication, which splits into spelled out (*coding*) and inferred (*expression*) meaning. Based on Watzlawick et al. [WBJ74], the spelled out communication content [McL64] can be regarded as coding, either in *digital* or *analogical* (sic!) form. The digital code roughly refers to the actual meaning of the transmitted information in a symbolic system (e.g., the writing “the sky is blue”), while the analogical code refers to



Figure 2.3: Conceptual framework of communication analysis systems. It consists of the four main dimensions *Input: Data and Information*, *Processing and Models*, *Visual Interface*, and *Knowledge Generation*, which allow for systematical analysis of such systems. Note that the graphic highlights some of the most relevant aspects for individual properties, which can be used for a simplified comparison. However, the properties themselves are multifaceted and, for a detailed analysis, should be discussed more nuanced and in more detail than indicated by these examples.

how something is communicated, including cues (e.g., biosignals, like winking, or emoticons). Analogical analysis is rare (e.g., message sentiment [HC14b; EGA⁺16]), partly due to the information loss in digital transmissions.

Expression — Similar, but orthogonal to it is the *expression*, which describes the intended or inferred information extracted from the content. It can be *explicit* factual information (the fact that the sky is blue) or information *implicitly* contained and which must be inferred, for example, from the semantics (or character [McL64] of the message). For example, the sky is blue - “let’s go hiking now”. Code and expression together allow classifying systems by their capability to leverage both digital (e.g., actual content information) and also analogical codes (e.g., inferred as sentiment analysis) while judging support for explicit (e.g., keyword-based search like [IBM20]) and also implicit (e.g., named entity recognition like [ESG⁺17]) content. Most approaches consider the explicit level, and several, especially text-based ones, also the implicit level.

Communication Participants

Of central importance are the participants in a communication. The *scale* of the communication is determined by its audience. In correlation with the *Context*, this determines different modes like narrow-casting (few, restricted participants), broad-casting (large audience), or targeting (specific participants), in turn influencing (or being influenced by) the communication *medium*). In communication research [McL64; WBJ74] these aspects are usually considered part of the context (see below).

Parties — According to the domain experts, the involved types of the parties can be another participant (with oneself as a special case) or different forms of groups (homogeneous groups or heterogeneous groups with subgroups) and differ between the sender and receiver sides. Therefore, we propose to structure the approaches based on their support to analyze the communication between a source and a target in a 3x3 matrix (individual, group, nested groups), e.g., individual to individual is encoded as the symbol  (e.g., no group support whatsoever [CSL⁺16]). Counterintuitively, the matrix might not always be symmetric.

Properties — The properties of these participant(s) can be manifold. One possible classification can describe them along their *capabilities* and their *experience* (knowledge of context).

Power Relationship — The (power) relations between the parties have a strong influence, with differences in *push* and *pull*. A possible classification [WBJ74] distinguishes between *symmetrical* (equal grounds) or *complementary* (dependence) relations, going beyond mere directionality. However, the relationship is rarely considered explicitly in existing research (e.g., [CLS⁺12] analyses the changing relations inside a group during an information diffusion process).

Context

The context of communication is essential [SW49; McL64; WBJ74], because it strongly affects the implicit interpretation among participants. We focus on the context of the external environment (*confidentiality, measurement, and medium*), and of the message (*timeframe*).

Confidentiality — The confidentiality of the communication channel can strongly influence the communication coding, for example through aversion or code-words (see also *factors* in Section 2.4.4 and our work [FHJ⁺22] on human factors).

Measurement Problem — Closely related is the measurement problem, where the analysis interferes and influences the communication coding and expression simply due to its (possible) presence, as indicated by the domain experts. The communication is affected by the participants' awareness, so they might adapt their behavior, use coded language, are less honest, implicit, or communicate not at all [FHJ⁺22]. This also concerns trust and reliability, both for parties and analysts [Cor19; FHJ⁺22]. We categorize this aspect into a quadrant, between (expected) knowledge about the analysis, which can be known or unknown to be true.

Timeframe — The timeframe when communication is occurring is highly relevant (e.g., for event correlation [SRG⁺19]). It can be described from the perspective of its duration and the activity during it [SFS⁺19].

Medium — The communication medium [McL64; Sch81] is partly covered (or mutually induced) by the participants and also the message type, coding, expression, and other contextual factors. Nevertheless, it deserves its own spot, in particular, due to media-typical characteristics and its relevance in research [McL64].

2.4.2 Processing and Models

After defining the data and information available, we study the particularities of processing and model creation from this information. We consider a technical perspective in visual data analysis, following the Golden Circle model by Simon Sinek [Sin09] to answer the *why?* (*Analytical Goal*), the *what?* (*Model Scope*), and the *how?* (*Processing*).

Analytical Goal

We start with *why* [Sin09], categorizing them by the aim of the analysis, which determines the analytical tasks to achieve it. We align our classification by the standard definition of analytical tasks in visual analytics [SGS18]. This includes the category *representation* for a fixed analysis task (present existing data), *confirmatory* analysis for a directed search (to validate a hypothesis about the data), and *exploratory* analysis for an undirected search (find interesting anomalies in the

data). Another goal in communication analysis involves *predictive* analysis (e.g., to analyze the future diffusion of information [WLY⁺14]), to draw conclusions, which can also be regarded as overarching all methodologies and is related to the knowledge extracted (see Section 2.4.4).

Scope

The scope answers the what?, determining the generic capabilities on the information. Other scopes are defined by their data (see Section 2.4.1) and knowledge generation (see Section 2.4.4) support.

Modality — The analysis modality is categorized into three [FSS⁺21] core aspects: *meta-data* (e.g., like time-series [KBK11]), *network* (e.g., social graphs [BHJ09]), and *content* (e.g., conversation order [ESKC18]).

Collection Type — The collection type is the logical composite to the *Measurement Problem*, defining how the data was acquired and its corresponding analysis implications. We propose to categorize it into a quadrant between targeting methodology and anonymity level (see the relevant part in Figure 2.3). The former can be either targeted (specific communication from a restricted set of users) or untargeted (unwarranted bulk collection). The latter can either be high (anonymized or pseudo-anonymized) or low (identifiable). Different configurations might pose particular challenges to the analysis model regarding aspects such as scalability and inference capability through class imbalance or uncertainty [FHJ⁺22]. For example, the targeted analysis of identifiable communication participants can focus on the actual exchange and leverage context and relationship information. The untargeted analysis of pseudo-anonymized communication instead often results in a search for the needle in the haystack and can rarely leverage background.

Processing

Due to existing heterogeneity, we focus on generic aspects: the *analysis* approach and *latency*, *scalability* as key performance indicators (KPI), and *data-mapping* as power.

Analysis — The employed techniques and algorithms often differ significantly between *offline* analysis and *online* analysis. Loading a dataset once would be considered the former type, while batch (e.g., updating data with changes [IBM20]) and, in particular, streaming approaches can be classified as the latter. Most approaches only cover offline analysis.

Latency — The latency is orthogonal to the analysis. Research [Pea11] indicates that latency in the communication can significantly affect it, as well as its analysis. The two primary options are (nearly) instantaneous communication, like in an active *live* (**L**) chat (e.g., live monitoring and analysis [Pal20]) or *delayed* (**D**)

) communication, such as e-mail or as a document. Differentiation into these two groups [Pea11] is often enough for most differences in reaction and behavior, although the latency can play a role (e.g., answering under time pressure).

Scalability — The scalability of a KPI can be defined on two levels: First, on the *data-ingress* level, which defines the amount a system can import, analyze, and visualize initially. The second aspect is the scalability on the *search* and analysis side, for example, during exploratory analysis. For example, how many results can be shown simultaneously? We roughly categorize both aspects into few (less than ten, I), medium (order of hundredths to thousands, II), and huge (more than 10k, III).

Data-Mapping — Supporting data mapping increases the analytical power of the systems. Supporting a flexible import system that allows mapping properties in contrast to a fixed data format is extremely important to the domain experts and often aligns with support for merging different data sources. For example, many systems cannot load multiple datasets and combine fields like usernames but only consider a single dataset (e.g., e-mails) in a fixed format.

2.4.3 Visual Interface

While there can be many design principles involved [CSJ⁺18], we describe the visual interface abstractly [KAF⁺08], focusing on three interrelated concepts: *representation* for the visualization, the techniques employed in *interaction*, and the synthesis of both through *refinement*.

Representation

The central aspect of visualization systems is their representations.

Method — We follow the established nomenclature of visualization techniques [KAF⁺08]. However, we only chose those common in communication analysis: *node-link-based* (e.g., [BM98]), *timeline-based* (e.g., [GS14]), and *matrix-based* (e.g., [FAS⁺20]). Other (e.g., chord diagrams [EGA⁺16]) techniques are grouped, while we additionally highlight *multiple-paradigm* (e.g., timeline, graph, and text [FZCQ17]) approaches.

Pane — The different visualization methods can be employed in different visualization panes. We consider the three major ones, namely **2D**, **3D**, and **S3D** (stereoscopic 3D like VR or AR). For example, a communication network can be visualized as a node-link diagram in either way, and each choice may influence the interaction concepts.

Interaction

Interaction methods are of central importance in visual analytics.

Table 2.2: Legend for Communication analysis system classification (cont.) shown on the previous page. The table summarizes the approaches properties. The classification criteria are formed on a subset based on the conceptual framework we developed in Section 2.4 and which is shown in Figure 2.3. The selected approaches and most categories are also available on a dedicated website at communication-analysis.dbvis.de, also permanently archived with OSF [FDS*22a]. For details, see Section 2.5.3.

Generic Properties:

Approach supports ● a property or not ○, or support is only very limited ◐ (e.g., show associated data without analysis).
 †: Commercial systems / widely used in industry.

The following encodings have special symbols:

All Matrix Symbols: For an explicit labeling, see explanations in Figure 2.3.

Group Communication: The approach supports different types of group communication analysis, as encoded in a 3x3 matrix. The rows (source) and columns (target) specify individual, group, and nested groups, respectively. For example, encodes that the approach supports only the analysis of communication between individuals and, additionally, of individual to groups. Groups as source are not supported, and nested groups not at all.

Latency: Support of live or a delayed analysis a posteriori.
Scalability: Tens (I), hundredths (II), or millions (III) of cases.

Evaluation: Case study , technique comparison , and expert study .

Time Dimensionality: Support for different time dimensionalities encoded in a 2x3 matrix. Rows form the knowledge basis (either past or present), and the columns specify if knowledge about the past, present, or future is inferred. For example, encodes that only past knowledge is used to learn about the past, e.g., to find a historic pattern. The encoding instead indicates usage of past and present data to predict future developments.

Predictive Power: The approach's predictive power encodes in a 2x3 matrix the type of output (explanation or transition, as rows) in relation to the time dimensionalities past, present, and future, as columns. For example, encodes an explanation about the past, while can also give explanations about the future. Instead, the encoding indicates that the approach supports both an explanation but also a transition function to explain past as well as classify present and predict future developments.

Operation Method — We classify the approaches based on their interaction method according to the classification developed by Yi et al. [YKSJ07], namely *Select*, *Explore*, *Reconfigure*, *Encode*, *Abstract/Elaborate*, *Filter*, and *Connect*. Some are extremely common, while others like *encode* depend on the capabilities.

Manipulation — The manipulation [KAF⁺08] of the elements can be either *direct*, for example, when interacting with data or visual objects. Alternatively, it can be *indirect*, for example, when modifying parameters. Most approaches support both.

Refinement

In addition to the interaction concept, other discerning factors are the particularities of the refinement, for which we differentiate [SSK⁺16] between the *goal* and the *strategy* to achieve it.

Goal — Two primary goals can be differentiated [SSSE20]: is the goal to tune an underlying *model* (e.g., for predicting communication behavior [FAS⁺20]) or the *data* (e.g., to select a fitting representation [LYW⁺16])?

Strategy — The refinement strategy [SSK⁺16] might vary: does it follow an *iterative* (e.g., improvements through continuous interactions [EGA⁺16]) or *progressive* (e.g., incrementally discovering events [KBK11]) strategy?

2.4.4 Knowledge Generation

Knowledge, generated and learned, is the ultimate analysis goal. We propose three subcategories: *output* to conceptualize the direct outcome, *knowledge gain* to cover the power of the outcome, and *verification approach* to consider implications and evaluations.

Output

Based on the classification of Spinner et al. [SSSE20], we propose two distinct categories for the learned knowledge type: An *explanation* can consist, for example, of numerical (e.g., graph algorithms [BHJ09]), textual (e.g., presented text [SRG⁺19]), or graphical representations (e.g., visual network representations [BN11]). It represents knowledge but in a factual representation that is not easily transferable and can be regarded as a (final) result of the existing data and is intended for humans. In contrast to this, another type of result can be a *transition function*, which is closer to an actual model, one example being the analyst's mental model. Another type is a machine model, for example, a trained, applicable classifier (e.g., diffusion model [WLY⁺14] or neural communication prediction model [FAS⁺20]) that encapsulates learned knowledge.

Knowledge Gain

As a final step in the learning process, the question arises which knowledge [GRF08] is actually gained and how powerful the process is.

Time Dimensionality — The time dimensionality describes the relationship between data and knowledge generation. A 2x3 matrix shows the possible combinations of data basis and prediction type, each with the entries past, present, and future. For example the symbol $\begin{bmatrix} \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \end{bmatrix}$ (like [HSCW13]) represents a system that can use past data and predict past data, for example for a search. Then, the encoding $\begin{bmatrix} \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \end{bmatrix}$ (like [LWY+20]) would reflect an article analysis and prediction system which has been trained on past data to analyze a text, either an existing one or one on the fly in the present and future. Another example for a future prediction is a model that forecasts communication activity based on past events, encoded as $\begin{bmatrix} \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \end{bmatrix}$. Note that by causality, the future is excluded.

Predictive Power — A second consideration describes the predictive power of the knowledge generated, which is represented as a 2x3 matrix, where the result (explanation or transition function) and the time are combined. For example the encoding $\begin{bmatrix} \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \end{bmatrix}$ reflects a system that can explain (i.e., show) past events (virtually all systems). The encoding $\begin{bmatrix} \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \end{bmatrix}$ represents one that can provide factual information for future events (e.g., information cascade prediction [LWY+20], inaccessible internal model). More powerful is a controllable model which can explain and predict (e.g., opinion diffusion [WLY+14]), encoded as $\begin{bmatrix} \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \end{bmatrix}$.

Domain-specific Aspects — Here, specific options depend on the analysis tasks. As discussed above, these are out of scope here; however, we could imagine this as future work (see Section 2.5).

Verification Approach

The presentation, as well as automated analysis of knowledge, raises a plethora of ethical as well as technical questions.

Factors — Visual analytics is well suited to address factors like confidence, trust and privacy, and consider aspects like fairness and accountability [FHJ+22; Cor19]. For example, probability scores could be used to estimate the results' confidence stemming from automatic processes (e.g., visual confidence scores [FAS+20]) and visualize it to the expert. Other examples include analysis log files, integrity protection, traceability, and verify-ability, as well as a provenance history. The lack of such certainty measures might exclude systems from sensitive areas. While essential, as shown by Correll [Cor19], many approaches are oblivious. For a more detailed discussion of the human factors in communication analysis, the ethical dilemmas, and design considerations for communication analysis systems, we refer to our companion work [FHJ+22].

Evaluation — To evaluate approaches, several options are possible: Either *examples* 📌 or a case study (e.g., describing an application [HFM07]). A second option is a *comparison* ↔ with existing approaches through feature comparison (e.g., comparing the most relevant tasks [FAS⁺20]). A third option would be a *qualitative* interview 👤 (e.g., interviewing eight domain experts [FSS⁺21]) or a *quantitative* user study 👤 .

2.5 Discussion and Future Work

We have defined four main dimensions containing over fifty different properties, providing a conceptual framework for interactive communication analysis systems employing visual analytics principles. This section discusses the main findings and lessons learned before reflecting on the difficulties in creating a conceptual framework for communication analysis. In particular, we evaluate the framework through two additional domain experts and discuss the potential implications and opportunities for future research while highlighting the shortcomings.

We imagine several **applications for this framework**: To provide a state-of-the-art overview of the current techniques, laying the foundations for a more detailed survey. Further, structuring the research field and providing a common language for the community while supporting comparison between approaches for practitioners and developers alike.

2.5.1 Expert Validation

To evaluate the relevance and validity of our conceptual framework, we recruited two additional domain experts for assessment. Both are analytical system developers and act as consultants to domain experts in intelligence, working in law enforcement. As such, they are highly familiar with the requirements and needs of the analysts as well as the technical implementation of the systems. According to them, the “framework covers many use cases in intelligence”, “misses no critical dimensionality”, and is well received as helpful to structure their collection of systems. As often “several systems are used in different stages of the investigation”, the framework is especially well suited to “assess the capabilities of presented approaches” for different tasks, to perform “market observation”, and the “basis for a detailed assessment scheme”. The framework “can support analysts in choosing” the correct system or sequence of systems. Currently, many investigation tasks “still require manual reading of all content” and analysis of relations, which is why more powerful and holistic approaches are needed. Systems that “hide the complexities”, work well, and are easily understandable can “increase the analysts’ performance significantly”. Due to the ethical and legal requirements in communication analy-

sis [FHJ⁺22], it becomes “increasingly relevant how explainability and information provenance” is handled and where a systematic assessment can be helpful.

2.5.2 Survey Findings and Research Opportunities

Our framework identifies gaps and research opportunities. Visual analytics is especially suited to support semi-automatic communication analysis [Her15; JFDK00]. The complexity, multi-modality, and ambiguity of communication are well-suited to interactively combine domain knowledge and computing. Concepts like interactive learning (e.g., [FAS⁺20]) allow refining models, while uncertainty awareness enables automatic judgments (e.g., [SFS⁺19]), fostering user trust and identifying bias [FHJ⁺22].

Findings — To apply the framework, we have taken the 41 selected approaches (see Section 2.3) and coded them according to our conceptual framework. Based on the results in Table 2.2, we can discover several interesting aspects: For one, studying the **data type**, while the analysis of text data seems mostly universal across representation methods, this is not the case for the other data types. Somewhat unsurprisingly, when network data is included, the visualization is often node-link-based or multi-paradigm, while for time-series, it is either timeline or multi-paradigm. Given the scope of the survey, the lack of audio and image is not surprising. Given that all the approaches belong to the category of visual analytics, it is also unsurprising that virtually all support representative, confirmatory and exploratory analysis as their **analysis methodology**, their **operation methods** covered most options, and their **explanation** is at least always graphical.

More interesting, however, are the **differences** and the **research opportunities** we can conclude from their discrepancies, which we highlight in the following for each category (see Figure 2.2), **linking back to visualization systems** as well as **communication research**; aspects we consider of particular relevance are highlighted with a star: ★.

I.1 Analysis of the Meaning and Analogical Code

Only a subset of visualization approaches analyzes the implicit meaning of the communication (e.g., [HC14b; ESG⁺17; FZCQ17]). However, almost none analyze the **analogical code** of the communication.

Implication: The analogical code can contain important cues which might support the analysis of the content and provide supportive information about the relationships inside the network, which makes it especially relevant to consider [WBJ74]. Leveraging it can lead to a richer and more-complete analysis, while it can support resolving contradictions and ambiguities [ESKC18].

I.2 Include Power Relations

Again, only a few (e.g., [CLS⁺12; HC14b; EGA⁺16]) approach considers the power relations between the participants, which can similarly influence the communication semantics, meaning, and modalities.

Implication: Power relations between participants [Sch81; WBJ74] might influence content aspects like the choice of words, formality, use of irony, meaning, or meta-data aspects like dynamics [SFS⁺19], timestamps, or message count. Results can be used and considered in context with the content analysis.

I.3 Dynamic Analysis

While some might consider this a technical problem, the development of systems that support the dynamic analysis of communication data and batch/stream approaches (e.g., [IBM20]) sets considerable hurdles to established analysis and visualization methods, which makes it an interesting research problem.

Implication: Exploring how new data and updated results can be integrated [CLS⁺12], how fluctuating analysis can be stabilized, and how changed predictions [LYW⁺16] are communicated offers more effective ways for visual communication.

Thesis Contribution: As part of Chapter 8, we propose a framework that includes batch and stream analysis of communication data.

I.4 Research the Measurement Problem

The measurement influence [FHJ⁺22; FSS⁺21] is rarely explored.

Implication: Being aware of the *measurement problem* and explore mitigations [SV19] can strengthen user trust, while avoiding missed or erroneous results (e.g., due to codewords) [WBJ74].

Thesis Contribution: As part of Chapter 3, we discuss privacy and ethical issues regarding the measurement problem as well as discuss potential solutions for the analysis. In Chapter 8, we propose a technique that can mitigate some of the posed issues, which we also discuss alongside a investigative journalism case study.

II.1 Multi-Environment Inclusion

Many approaches lack support for data mapping and multiple data sources (e.g., [IBM20]), requiring preprocessed data.

Implication: Automating the merging of heterogeneous data sources [KAF⁺08] with few or no user input reduces the amount of manual preprocessing or knowledge transfer required, makes leveraging multiple data sources simultaneously less complicated [FZCQ17], while exploring optimal interface strategies.

Thesis Contribution: As part of Chapter 8, we present a system architecture that allows for an holistic and multimodal integration of communication data.

★ II.2 Analyze Group Communication

Only a few approaches support the analysis of group communication (e.g., [FM08; EGA⁺16]), and almost none support nested groups.

Implication: New and more detailed knowledge can be drawn on how groups operate [FM08] and information diffuses [Lea51] within, in particular, because much communication actually happens inside or between groups, which can involve specific particularities [Bav50].

Thesis Contribution: As part of Chapter 5, we contribute methods for the identification of communicating groups.

III.1 Visually Interactive Model Analysis

Virtually all approaches use the 2D pane for visualizations, and many automations focus on filtering instead of model tuning.

Implication: Leveraging visual data analysis techniques [KAF⁺08; YKSJ07; Her15] and explore unused approaches like VR for improving the analysis process [CSJ⁺18; GKL⁺13], focusing on the model [SSSE20] instead of only data selection may allow for the higher level conclusions, supporting the knowledge generation [SSS⁺14].

Thesis Contribution: As part of of thesis, we propose several visual analytics techniques: we investigate its theoretical potential in Chapter 3, and present HYPER-MATRIX (Chapter 5), COMMAID (Chapter 7), Conversational Dynamics (Chapter 6), as well as MULTI-CASE (Chapter 8) as prototypical solutions.

★ IV.1 Model / Transfer Function / Knowledge Gain

Few approaches contain an actual, powerful machine models (e.g, [WLY⁺14; ESKC18; Nui20; FAS⁺20; FSS⁺21]) to analyze communication.

Implication: Using such models can potentially support the analysis [GRF08; CSJ⁺18], through measures as active learning [CSL⁺16; NNW15], intelligent filtering [FSS⁺21], or confidence-based predictions [SSK⁺16]. Transfer Functions allow for a more universal machine learning, applying knowledge to new problems, increasing the predictive power. This reduces manual work while increasing analytical capabilities.

Thesis Contribution: While all the previously mentioned approaches (Conversational Dynamics in Chapter 6, CommaID in Chapter 7, and MULTI-CASE in Chapter 8) also support the refinement of internal models, the model retraining and adaption process is investigated in particular detail as part of HYPER-MATRIX in Chapter 5.

★ IV.2 Confidence, Trust, and Privacy

These factors are insufficiently considered in the majority of approaches, leading to a black-box analysis. Instead, one could include confidence estimates (e.g., [FAS⁺20]), logs, provenance (e.g., [FSS⁺21]), data minimization, or other concepts.

Implication: Several applications have strong requirements for confidence and trust [SSK⁺16], provenance [SSSE20; FSS⁺21], and privacy [FHJ⁺22]. Exploring how these can be fulfilled [FHJ⁺22] without limiting the analysis can replace manual analysis by automated systems.

Thesis Contribution: We discuss the confidence, trust, and privacy factors as well as the ethical implications in detail in [Chapter 3](#). The lessons learned are also reflected in the design of MULTI-CASE presented in [Chapter 8](#).

IV.3 Guidelines and Quantitative User Studies

While several approaches include case studies and (qualitative) expert interviews, almost none make actual comparisons with related approaches or conduct quantitative user studies.

Implication: Case studies and qualitative expert interviews are not always comparable or conducted to the same standards [BISM14]. While we do not doubt the systems work well as advertised, for reproducible comparisons between approaches, quantitative studies are required and evaluations along design guidelines [CSJ⁺18], providing a more objective overview.

Thesis Contribution: While quantitative user studies in this domain remain extremely difficult, we present guidelines and points to consider as part of the design process in [Chapter 3](#) as well as [Chapter 7](#). Further, we compare our systems HYPER-MATRIX ([Chapter 5](#) and [Chapter 4](#)) and MULTI-CASE ([Chapter 8](#)) with related approaches or according to a capability assessment.

★ O.1 Holistic Approaches

Few approaches perform a holistic analysis (e.g., [FSS⁺21]) by considering multiple analysis aspects in context, covering all modalities.

Implication: A holistic perspective [v14] can increase the analytical capabilities [BISM14; FSS⁺21], supporting cross-matches beyond analysis boundaries [FHJ⁺22], while reducing manual and mental load.

Thesis Contribution: As part of this thesis, we present two holistic approaches: COMMAID ([Chapter 7](#)) and MULTI-CASE ([Chapter 8](#)).

★ O.2 Context / Analysis Reference Window

Similar to the holistic analysis, a specific focus on the communication context

in reference to each other should be explored further for both inter- and intra-modality analysis.

Implication: Few approaches consider other modalities or external factors to explain particularities. For example, a break in a communication sequence might appear as a gap, but when combined with location information (e.g., same building) might indicate that the participants might have met for lunch and continued their conversation offline. In summary, the correct interpretation of communication is extremely context-dependent [CSL⁺16; CSJ⁺18], with different applicability of analysis methods. Analyzing references and clues can improve the determination of the highly variable context [Mes09] for choosing appropriate analysis methodologies.

Thesis Contribution: We address the question of inter- and intra-modal analysis to various degrees in different chapters: We investigate the identification of missing links through common context and shared similarities as part of HYPER-MATRIX in Chapter 5, the identification alongside common communication meta-data patterns as part of Conversational Dynamics in Chapter 6, and a more general context in both COMMAID in Chapter 7 as well as MULTI-CASE in Chapter 8.

Implications — All the previously described opportunities offer potential improvements for a more complete analysis of communication. Fusing together multiple methods can lead to a richer and more complete analysis, potentially resolving contradictions and ambiguities. Some opportunities primarily support existing analysis steps (e.g., I.2), while others provide new areas (e.g., II.2, O.1). While the relevance of aspects might differ in any given analysis, our framework identifies areas that users - and developers of interactive analysis systems - can consider and leverage, depending on their analytical needs.

2.5.3 Limitations and Future Work

One problem with the described taxonomy is the basis it is designed upon, which can affect its **completeness**). A significant problem in this research area is that relevant approaches are rarely labeled as belonging to communication analysis. We initially thought about compiling a list of domain-specific keywords for selection (e.g., social network analysis, sentiment analysis, e-mail analysis, etc.). However, we found it highly likely that such a selection would be highly biased by our knowledge, which is why we decided to take a wider approach. However, even with care, it is inadvertently likely we missed individual approaches, not least by restricting the target journals and interviewing domain experts from (criminal) intelligence, although related domains like investigative journalism have similar requirements [FHJ⁺22]. Also, it

could happen that a few approaches fell through our automatic or manual search pattern (see Section 2.3). However, due to the restrictions discussed in Section 2.3, we do not claim overall completeness. The survey forms one of three pillars for our goal of constructing the framework, and together with the other two, we are confident the majority of cases can be described within our framework. Nevertheless, to address the issue of missing approaches, we created an accompanying **survey website** available at <https://communication-analysis.dbvis.de>, which lists the approaches we considered and also allows readers to submit methods missing methods. This website is also permanently archived with an OSF repo [FDS⁺22a].

Another possible limitation concerns the **orthogonality** of the framework itself. Due to the complexity and heterogeneity of the area, it contains some overlaps. As there is a need to balance the trade-off's accuracy, usability, and relevance, we think it is challenging to create a wholly consistent yet easily usable taxonomy. The choices we made for selecting the categories are often based on the literature and justified when required. However, given sparse taxonomy and non-standardized vocabulary, some groupings and namings could arguably have been chosen differently with the same validity. To advance research in this area, however, we decided to propose our framework as a first possible draft and one step towards a universally accepted framework. We, therefore, invite the research community to give feedback to stimulate the scientific discussion and extend the framework through input from diverse research communities. As part of this process, the individual, multi-faceted aspects can be formalized in more detail.

Another aspect is the extension of the framework to **non-human communication**. Several aspects could be applied to communication in general, for example, machine to machine. Indeed, nothing in the framework is specifically tailored to a human communicator. However, human communication is often more nuanced than machine communication, making parts of the framework less relevant, while other features (e.g., structuring, exchange content scope) might be missing so far.

2.6 Conclusion

In the last decades, communication analysis has experienced a shift away from manual analysis to computer-aided or even highly automated approaches. However, to the same extent as automation levels increased, the analysis itself has often become more specialized, moving away from an overarching exploration. This trend is in contradiction to traditional communication research, which stresses the importance of a holistic approach to capture the full meaning and context of communication. As a result, many modern digital communication analysis systems are highly adapted to a narrow range of tasks, often either in the area of content or in network analysis. While this might be perfectly sufficient and suitable for

their intended use, such an isolated analysis can sometimes lead to a less effective exploration and lead to incomplete or biased results. Using separate approaches requires more manual work, often complicates analysis tasks, can introduce domain discontinuities, and increase the struggle domain experts face when trying to integrate their knowledge. Further, an isolated analysis may not be sufficient to capture the full available information and can make the automatic as well as the manual detection of cross-matches more difficult. The development of more holistic and advanced approaches for automated communication analysis systems is hindered by the lack of a clear framework and the absence of a common language that combines technical aspects and traditional communication research.

We address this challenge by developing and formalizing a design space for digital communication analysis systems based on the existing tool landscape and communication research while making a case for how visual analytics principles can be employed for a more holistic approach. By systematically discussing and structuring the different analysis areas and aspects of the design space, we arrive at a conceptual framework to provide an overview and assess the maturity of communication analysis systems. As part of a state-of-the-art survey, we have also categorized a large set of existing approaches using our framework.

By bridging the gap in the formalization of digital communication analysis systems by describing a design space for communication analysis, we aim to provide researchers with a common language, provide guidelines for building and assessing the maturity of such approaches, as well as point out gaps in the literature which offer exciting research opportunities. Further, the formalization acts as the conceptual basis for the remainder of this thesis. The results are widely applicable in a variety of domains that are concerned with communication analysis like national security, the digital humanities, or business intelligence, both from a theoretical point of view as well as for the development of more powerful communication analysis systems.

Ultimately, saying that you don't care about privacy because you have nothing to hide is no different from saying you don't care about freedom of speech because you have nothing to say.

— Edward J. Snowden, *Whistleblower*

3

Ethical Considerations: Potentials, Limits, and Implications

With the increasing abundance of data that can be harnessed by analyzing human communication, manual management seems increasingly unrealistic while posing new challenges regarding automated behavior. Particularly interesting cases are formed by the intelligence analysis of communications data in areas such as investigative journalism, criminal intelligence, and law, as they must take into account the often highly sensitive properties of the underlying operations and data. At the same time, these are areas where increasingly automated, sophisticated approaches and tailored systems can be particularly useful and relevant, especially in terms of Big Data manageability. However, the shifting of responsibilities also poses dangers. In addition to privacy concerns, these dangers relate to uncertain or poor data quality, leading to discrimination and potentially misleading insights. Other problems relate to a lack of transparency and traceability, making it difficult to accurately identify problems and determine appropriate remedial strategies. Enabling human sense- and decision-making as a joint agency, for example, through the use of visual analytics, can be key for designing and operating meaningful interactive communication analysis systems that consider these ethical challenges.

In the following chapter, we investigate and evaluate opportunities and risks involved in using Visual analytics approaches for communication analysis in intelligence applications in particular. We consider the apparent challenges from an interdisciplinary viewpoint, reconciling ethics, science and technology, and com-

Contents

3.1	The Ethical Dimension	65
3.2	Ethical Challenges for Communication Analysis	67
3.3	Scenario Analysis	69
3.4	Mitigation Techniques through Visual Analytics	75
3.5	Conclusion	84

puter science perspectives. At first, we discuss common technological systems used in communication analysis, with a special focus on intelligence analysis in criminal investigations, further elaborating on the domain-specific ethical implications, tensions, and risks involved. We then make the case of how tailored Visual Analytics approaches may reduce and mitigate the described problems, both theoretically and through practical examples. The latter parts of this dissertation (Chapter 5 onwards) will then discuss actual techniques and implementations of these approaches. In the following, we argue how offering interactive analysis capabilities and what-if explorations while facilitating guidance, provenance generation, and bias awareness (through nudges, for example) can improve analysts' understanding of their data, increasing trustworthiness and accountability while generating knowledge. We show that finding Visual Analytics design solutions for ethical issues is not a mere optimization task with an ideal final solution. Design solutions for specific ethical problems (e.g., privacy) often trigger new ethical issues (e.g., accountability) in other areas. Balancing out and negotiating these trade-offs, as we argue, has to be an integral aspect of the system design process from the outset. Finally, we identify existing gaps and highlight research opportunities, further describing how our results can be transferred to other domains. With this contribution, we aim to inform and build the ethically-aware basis for communication analysis in this dissertation and in the intelligence domain more generally.

This chapter is based on the publication [FHJ⁺22] and major parts of the following sections have appeared in:

- [FHJ⁺22]: **Maximilian T. Fischer**, Simon D. Hirsbrunner, Wolfgang Jentner, Matthias Miller, Daniel A. Keim, and Paula Helm. "Promoting Ethical Awareness in Communication Analysis: Investigating Potentials and Limits of Visual Analytics for Intelligence Applications". In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, 2022, pp. 877–889. ISBN: 978-1-4503-9352-2. DOI: [10.1145/3531146.3533151](https://doi.org/10.1145/3531146.3533151).

For a statement of the scientific contributions, as well as the division of responsibilities and work in this publication, please refer to Chapters 1.2 (p. 22ff) and 1.3 (p. 26ff), respectively.

3.1 The Ethical Dimension

The intelligence community is one of the most prevailing domains for sophisticated big data analyses, namely in **criminal investigations**, lawsuits, matters of national and international security and in **investigative journalism**. Specialized systems are used, for example, by the National Security Agency (NSA) as part of global spying operations [MJ17] or law enforcement against organized crime [DFH⁺21], by lawyers for analyzing case-relevant documents [ATPL16], but also by journalists working on [WYB18] large data leaks such as the Panama Papers. During these operations, large amounts of communication data, like e-mails, chats, posts, or calls are collected, along with associated documents (e.g., attachments) and meta-data like timestamps, locations, and contact networks. In our research context, these domains present a particularly interesting case, as they should consider the often highly sensitive and private character of the underlying operations and data with particular caution. There is no doubt that untargeted mass collection of communication in the name of national security is privacy-invasive and thus highly controversial [Mac18], at least for western civilizations. The belief in the separation of the private and public life originates in Ancient Greece during the time of Pericles. There—as delivered by Aristotle—the concepts of *oikos* (meaning family or family property) and *agora* (meaning central market) were conceived, referring to the private and public aspects of life [Are60]. However, many ethical challenges remain even relevant for morally more accepted cases like specifically targeted analysis of confiscated organized crime equipment or even practices considered essential to democratic culture and particularly valuable, like data journalism.

For example, **privacy** issues relating to the separation of irrelevant data, its secure handling, analysis, and deletion have to be considered [Aus12; Amo14]. Poor data quality, unreliable methods, or biased algorithms may lead to misleading insights [Ame16], bearing the risk of overlooking critical information, or worse, contribute to discriminatory practices against people of colour [Ben19], cement existing social inequalities [Eub18], and may even result in false accusations [ONe16]. Moreover, a lack of **transparency** can make it impossible to defend oneself against such accusations [Sel17] when the systems supporting (or making) such decisions are considered reliable, but actually do not (consistently) provide complete chains of evidence [Kit17]. Unfortunately, being focused on efficiency and quick results, not all actors consider the arising ethical challenges in this field with a long-term perspective in mind, and even those who try, may be limited by their technical approach and implementation difficulties, for example, through black-box machine learning models or an incomplete understanding of the considerations involved in deriving the result [Zar16].

The concept of **accountability** in computer science [Nis94] stresses the need to handle and answer to the harms and risks that can be caused by technology. While this concept has been around for decades already, concrete ways of handling accountability in digital analytical systems have remained vague. Yet, progress has been made in interpretability of machine learning [Mol19; Rud19]. Whereas the coming into effect of the EU General Data Protection Regulation (GDPR) led to more awareness on this topic and its implementation [HvB18], its legal effects on the fields under consideration are limited due to exception clauses in article 2 (law enforcement) and article 85 (journalism) [Eur16]. Even in those areas, however, the need to consider these topics carefully is increasingly prevailing.

Working together closely with criminal investigators from various institutions, we know of a growing **awareness** of these difficulties, also on the part of the analysts themselves. Concerns about the trustworthiness of their analytics systems and the ethical considerations involved have been expressed. This concern is also reflected in digital communication analytics in general [vP18], where the need for more detailed analysis was identified. To date, ethical concerns related to automated communications analysis have been described mainly from either a strictly sociological and/or ethical perspective [HH21] or in the context of technical capabilities [vP18; FL20; FDS⁺22b]. Much less work, by contrast, addresses the complex techno-ethical tensions and dilemmas that arise in the messy gray areas of socio-technical feasibility, given the limits and consequences induced by the alleged solutions. Given the lack of overarching work, in this chapter we examine ethical considerations in communications analysis for intelligence applications in more detail and propose possible mitigation techniques, which we discuss critically with regard to ethical concerns. Unlike most previous research, we thus bridge ethical considerations, sociological science & technology studies, and a computer science perspective.

We study concrete design approaches and solutions and by analyzing the **interfacing** problem from an **interdisciplinary** perspective, we can critically reflect on the opportunities and challenges involved [Lip17]. We argue that designing solutions is not a mere optimization task, but balancing out and negotiating the trade-offs has to become an integral aspect of the design process at the very outset. We further claim that — in light of recent requirements for human oversight [Eur21a; Eur21b] — an application design based on visual analytics principles is uniquely suited for such a task, by interactively combining human sense- and decision-making, controlled through a frequent feedback loop.

In this chapter, we aim to promote ethical awareness of digital communication analysis by turning our focus to the interface, investigating the potentials and limitations of visual analysis for intelligence applications, contributing:

- A detailed **discussion** on the ethical frictions and tensions involved in intelligence applications, followed by a **scenario**-based stakeholder analysis of actors and their roles.
- A **critical reflection** of visual analytics design solutions fostering ethical awareness in communication analysis and the involved trade-offs as an integral part of the interface design process.

With this contribution, we aim to inform more ethically-aware approaches to communication analysis in intelligence operations using visual analytics principles.

3.2 Ethical Challenges for Communication Analysis

The field of AI ethics is a novel field. It concerns itself with the ethical implications and consequences of the increasing automation and datafication, both facilitating the discussion and responding to the various concerns raised. As such, it has seen substantial growth in the last few years. Given the considerable potential to misuse novel technologies in police intelligence [Sel17], they have received particular scrutiny when becoming public.

Both the academic community specializing in ethics as well as civil society actors heavily criticize the sometimes ill-advised deployment of such technologies. This is particularly well illustrated by the strong reaction against predictive policing technologies (PPTs) [Ame16]. In the following section, we will discuss several pertinent challenges (**C1–C6**) in this growing debate. Thereby, we focus in particular on those issues with significant implications for communications analysis.

C1. Discriminatory Bias – The replication of pre-existing stereotypes and exacerbation of community discrimination (e.g., women, people of color, transgender) by biased algorithms is a pressing challenge [Ben19; Cos20; Nob18; LI16]. Bias in machine learning can originate from human actors (e.g., suspicion just based on ethnicity) or be introduced through processes (e.g., the mirroring of stereotypes embedded in the training data). The latter can be observed in facial recognition systems [GF16] and pre-trained language models [KMN20]. To alleviate such statistical discrimination in AI models, one can develop fairness principles and metrics to assess and try to mitigate such issues [MMS⁺21; BHN17]. However, these measures may not suffice in communities already subject to inequitable conditions, necessitating an emphasis on equity prior to automation [Gre18].

C2. Privacy – Privacy remains a central challenge. Intelligence applications process copious amounts of data ingested from diverse, often confidential sources, including recorded phone calls, online messages, and social media. Most of this data pertains to unrelated third parties due to its sheer volume and heterogeneity,

thus raising significant ethical and legal concerns when stored and processed [MC16]. Although law enforcement may be exempt from many data protection and privacy regulations (most notably the GDPR [Eur16]), this does not absolve them of their social responsibility. It emphasizes the need for robust measures to prevent misuse of the collected data and clearly define, explain, and implement data access in a legal and technically secure manner.

C3. Opacity – Many current deep-learning systems remain opaque and can be considered black-box, raising many important questions [Pas15]. For this, one has to consider that even if such system *would* be interpretable and explainable to machine learning experts, typical analysts usually are no such experts. This makes it challenging for them to comprehend the system-generated output [AC18]. Further, other problems are typical for public-private partnerships: for example, trade secret protection for algorithms used in public domains can lead to severe difficulties in verifying the inner workings. This offers limited opportunities for analysts, judges, and the general public to challenge findings based on biased or inaccurate models, which sometimes can have dramatic consequences [LI16; ONe16; Nob18].

C4. Exaggerated Expectations – Another critique addresses the general misconception that algorithmic recommendations are indisputable because they are mathematical truths. Contrary to the black-box discourse, this discussion is ideological, emphasizing trustworthiness, not from a technical side. This concept of *mathwashing* should not be used to highlight the potential of innovative technologies, where software is framed as the single solution against any human errors or ill-intent [Joh17]. The portrayal of charismatic machines [Ame15; Ame19] triggers this fantasy. A prominent example of this disparity between promise and actual performance is served by predictive policing technologies [HH21] STS scholars underscore the importance of recognizing the subjectivity and intent embedded in systems, destroying the flawless machine's connotation due to collective decisions by diverse actors with potentially conflicting interests involved in its design and usage [Mac15; Kit17]. Communications analysis of intelligence data is far from neutral, considering the collective efforts and design decisions behind it [Git13]. The ability of people to question the validity of results is compromised when they believe in the neutrality of such tools when, in fact, they serve specific needs and interests. This concept is called *erasure of doubt* and can impede the application of previously learned and experienced trust [Amo19].

C5. Human-Machine-Configurations – The complexity of human-machine configurations has grown alongside advancements in automation [Suc07]. A key challenge lies in achieving both effective integration and establishing an appropriate level of automation. The automated analysis aims to aid investigators by filtering out seemingly irrelevant patterns and relieve them of repetitive or laborious tasks, allowing them to concentrate on more useful activities [Cor19]. However, when

analysts feel marginalized and displaced in their human experience through their machine counterparts, considering their combined work as a competition rather than assistance, this becomes problematic [KEL18]. The crux of the matter revolves around determining a sweet spot regarding the extent of automation, how it should be applied, and how it should complement the activities of the human investigator [SJB⁺17]. When and which alternative search terms or associated individuals should be suggested, what agency should the analyst retain, and how much and what type of contextual information should be displayed or concealed for privacy considerations? These are vital questions with no easy answer. Likewise, how much should interfaces be designed to prompt users to contemplate ethical issues, such as through nudging? Smooth and well-structured collaboration is a highly debated topic as it is seen as an essential safeguard against AI systems potentially undermining human autonomy [Eur21c] with many potentially harmful effects.

C6. Accountability – The ethical challenges discussed thus far also pose questions regarding accountability, which describes the readiness to assume responsibility for actions and decisions taken and extends not only to the users but also the software and the designers of its AI [AI 20; LOL⁺18]. A consensus needs to be found regarding the extent system users can and should be held accountable for consequential errors in the context of AI if the software fails to meet basic explainability and interpretability standards. Additionally, the obligations of the software provider to protect and safeguard against ethically questionable decisions (such as racially-biased categorization of suspects) and uses (like spying on third parties) need to be outlined and (legally) discussed.

3.3 Scenario Analysis

As a prerequisite for the following discussion, we first provide a overview of communication analysis and the common technological systems, before presenting the PEGASUS research project as a case study. We then construct a hypothetical scenario, from which we derive a map of the stakeholders in conflict, forming the basis for our proposition for mitigation.

3.3.1 Digital Analysis and Employed Technology

Digital communication analysis as a research field has no universally accepted definition, with different understandings in different domains. In this work, we follow our definition [FDS⁺22b], defined in Section 2.1, considering it to encompass the computer-mediated [FSS⁺21] analysis of meaningful digital [Sco09] information exchanges between humans [Pea11]. The analysis relates not only to the actual content (text, audio, or video), but also encompasses accompanying meta-data

as well as communication network structures. Existing communication analysis approaches rarely consider these aspects holistically [FSS⁺21], but primarily focus on individual aspects: Most commonly, these are textual analysis through fuzzy search (and increasingly natural language processing (NLP) [MS99] methods) as well as social network analysis [Sco17]. For example, in intelligence, one of the most commonly used systems [FAS⁺20] is IBM's i2 Analysts Notebook [IBM20], which has a strong focus on network analysis and information management but has, so far, lacked advanced textual analysis capabilities. However, competing solutions such as Nuix [Nui20], DataWalk [Dat20] and Palantir Gotham [Pal20] have been gaining ground [FSS⁺21]. Many are primarily large information management systems, using established algorithms (e.g., for centrality calculations in a network) and deterministic filters (e.g., keywords). Novel machine learning-based capabilities used for relevance scoring, person attribution, or facial matching are increasingly used in this context. The reliability of these models, however, the question of hidden bias, and the overall reproducibility (e.g., after updates), remain unclear.

In investigative journalism, tools like New/s/leak 2.0 [WYB18], as, for example, used by *Der SPIEGEL*, use models trained on public data like Wikipedia for discovering named entities in textual data (e.g., persons or company names). Similarly, the industry-standard spacy [Hon19] uses public corpora and increasingly open web information for model training. While this often results in increased accuracy, concerns about the reliability for less common languages or risks of manipulation (e.g., for datasets extracted from Wikipedia) remain valid.

3.3.2 The PEGASUS Research Project

For a case study on the requirements in intelligence, we specifically focus on the insights gathered through the work in the academic research project PEGASUS, funded by the Federal Ministry of Education and Research of Germany (BMBF). The project aims at improving big data analysis in the context of civil security, also considering the ethical challenges involved. The PEGASUS acronym — *not* to be confused with the unfortunately equally named PEGASUS spyware — stands for *Collection and analysis of heterogeneous Big Data by the police to fight organized crime structures*. Organized crime is a transnational and global form of crime, encompassing a broad spectrum of different areas, including human, drug, and arms trafficking, money laundering, smuggling operations, environmental, medical, cyber, and other white-collar crimes. According to Europol, in Europe alone, the number of criminal organizations under investigation is over 5000 (2017) [Eur21d], coming with high economic cost and a destabilizing effect on public security (through, for example, extortion, fraud, trafficking, or bodily harm). Organized crime can be characterized by its organized hierarchies (e.g., clans, mafia structures, shell

companies) and sophisticated criminal acts using modern technology, and their ability to adapt quickly to changing circumstances [Pao02]. For example, the COVID-19 pandemic has significantly affected organized crime, which was quick to adapt to new illegal avenues and *modi operandi* [Eur21d]. In conjunction, the seized data is increasing massively, overwhelming traditional (primarily manual) investigation methods. A significant share accumulates as intra- and inter-group communication and can be acquired, for example, when electronic devices are seized. However, the challenges faced are not unique to law enforcement; the goals are strikingly similar to tasks in fields such as investigative journalism and business intelligence, where information and the knowledge derived from it have become more important than natural resources [KN98]. Tackling the arising ethical issues is challenging because mitigation techniques incorporate numerous tensions and dilemmas that must be carefully weighed between the complex interplay of actively and passively involved stakeholders.

3.3.3 A Scenario in Police Intelligence Work

We construct a hypothetical scenario [Car99] of communication analysis within police intelligence work using current but non-visual analytics software that acts as a reference for our study of ethical challenges and emerging mitigation strategies. The scenario focuses on the challenges and practices of police officers and investigators as the main user community and points out other actors and stakeholders (highlighted as **SName**) in an exemplary way.

SMartin is a police officer at the organized crime unit of the federal police. He currently investigates the selling of fake COVID-19 vaccination passports by an alleged criminal organization named *The Medics*. The Medics offer the counterfeit certificates to their **Scustomers** via the Telegram messenger. Unknown to Martin yet, **SChris**, **SCarlos**, and **SEggert** are Medics members, also communicating with their colleagues and suppliers via group Telegram channels while using pseudonyms, sometimes coded language, and images. In their free time, they also communicate with several friends, including their girlfriends, **SSarah** and **SMarta**, who are unaware of their business. Martin's police unit gathers much information about The Medics using traditional investigative methods. This information leads to the identification of the suspect, Chris, who seems to be a low-level member of The Medics. On one evening, Chris is found with blank vaccination certificates during a traffic stop. He is arrested, and his phone is seized by investigator Martin, who aims at using the information on the phone to track down the individuals pulling the strings. After calling judge **SRobert** to get a search warrant, which is granted, he then searches Chris's unprotected phone, finds the Telegram communication, and extracts it. He recalls that his superior, **SDr. D**, asked him to try out the new

AutoCommAnalyzer software, which was recently purchased from the multinational company AI-Tech Corp. The software purchase was part of a strong push by the government to digitally optimize work processes at the police forces. Martin looks at the training notes by the head developer (SMolly), trying to remember how the machine learning-driven software — trained with texts by (SAlf) and (SBert) — is supposed to direct him to the relevant communication. The software presents him with the most frequent contacts, with Sarah on top. He reads through this communication, as the software has flagged several words like package and hospital, discovering some explicit images but finding that the flags refer to a delivery package and a hospital stay for a broken ankle. In a second chat, the AI highlighted several currency amounts, and manually reading through it, it becomes clear that expensive “stamps” have been sold. Luckily for him, many addresses and names are also included in the chat messages. Searching for all chats that talk about stamp selling, he also finds one with Carlos, including his last name, and one from a person called Big E, which includes an address. Using the nationwide register, he finds a person named Carlos, who used to be a roommate with Chris, and only one person named Eggert is living at the found address. After completing his analysis, he finishes his report and submits it for the trial. During the court proceedings weeks later, Martin is questioned by judge (SMuller) on his findings. Ultimately Chris, Carlos, and Eggert admit to their guilt and are sentenced for document falsification. The intelligence gathered points to other alleged criminal networks and informs other running investigations.

3.3.4 Conflicts of Interest between Different Stakeholders

A stakeholder analysis based on the previous scenario helps to identify, map and describe the different actors and their roles involved in the scenario (see Figure 3.1), with the interdependent individuals having potentially conflicting interests. We propose four main groups of stakeholders: *civil society*, *governmental authorities*, *software provider*, and *data subjects*. This categorization has to be understood as a heuristic with possible overlaps and without claims of being exhaustive.

Data Subjects — Following the concept of data subjects defined by the GDPR [Eur16], we consider the role of natural persons and their data ownership. Immediately apparent becomes the role of the **targets** (SChris). In many cases, a target is unknown, but one has a list of **suspected targets** (SCarlos), indicating a different degree of certainty. One issue of communication analysis, however, is that communication is not strictly separated and touches on many other data subjects. These can be as of yet **unidentified** persons (e.g., the customers), but also **third parties** (e.g., SSarah). With the use of machine learning, a fifth subgroup emerges, the **training data subjects** (SAlf and SBert), whose data is leveraged as part of

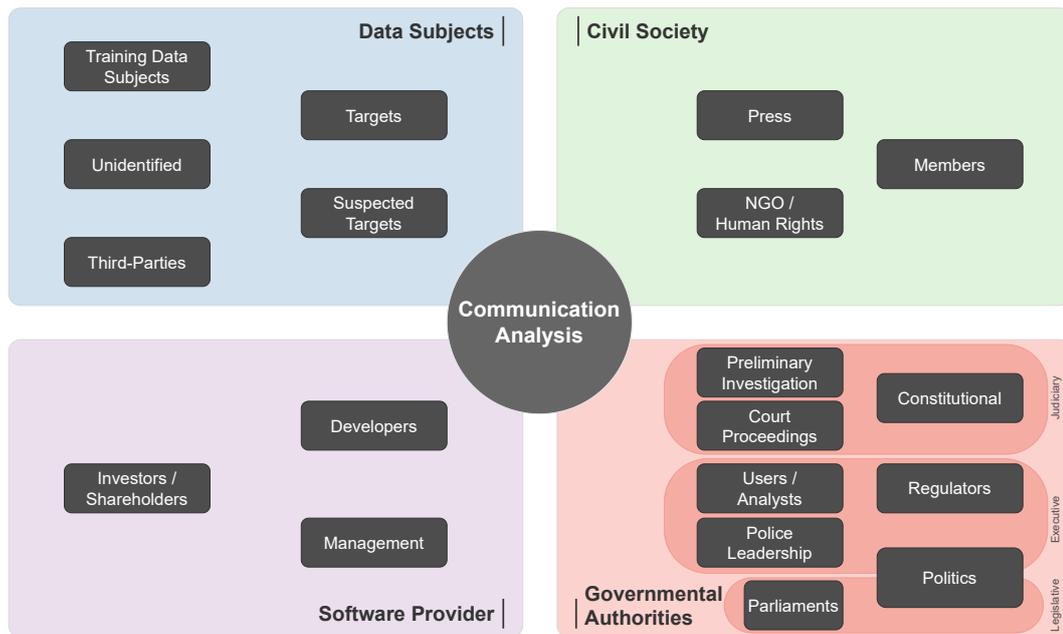


Figure 3.1: The stakeholders involved in communication analysis from the perspective of intelligence analysis, with conflicting interests giving rise to ethical dilemmas. We propose four main pillars of stakeholders: the *civil society*, the *governmental authorities*, the *software provider*, and the *data subjects*, each with its own subgroups of stakeholders, such as targets, developers, analysts, or NGOs.

training the weights in neural networks. Further conflicts of interest arise between uninvolved third parties who usually (and rightly) do not want to be involved in a privacy-invasive investigation, which can also apply to (unwitting) training data subjects. A delicate issue are privacy considerations in the face of imminent suspicion: while target subjects clearly do not want to be investigated either, the reasons and justifications here differ substantially to those of uninvolved third-parties.

Governmental Authorities — In this specific type of intelligence analysis, the opposite of the data subjects are governmental authorities, with their investigating bodies. Here, because this applies to our setting, we assume a democratic political system that follows a separation of executive, legislative and judicial powers. The investigating bodies are primarily part of the *executive*, with police **analysts (users)** (S Martin) conducting the investigation, overseen by **police leadership** (S Dr. D) and also controlled by **regulators** like data protection or compliance offices. The *judiciary*, however, also plays a controlling role during **preliminary investigations** (S Robert), allowing for specific measures, for example, by issuing a warrant. Later, it manages **court proceedings** (S Muller) and questions of legality can ultimately be decided on a **constitutional** level. The third power, *legislative*, is not directly involved in investigations but sets the boundary conditions for law enforcement through regulations, usually through **parliaments**. The area of **politics** employs a

ambiguous role in this case, influencing decisions but also constituting a part of the executive. Conflicting fields can arise between all government levels. For example, the top levels might put pressure on the bottom to produce results, promoting automated analysis for its efficiency. Analysts, in turn, may use legally questionable methods, the judiciary may be concerned about the failure of legal proceedings in such cases, and regulators may be concerned about established practices that run counter to the intentions of Parliament. A recently observed problem occurs when the relationship between the system implemented as an assistant and the sovereign analyst is reversed. This can lead to effects resembling defensive decision making, where police officers intentionally make suboptimal decisions by following the results of the machine "assistants" even when they disagree. This is mostly explained by pressure from "above" and the need to protect themselves from redress if something goes wrong [AAG19].

Software Provider — The software provider develops the tools officers use in their investigations. Here, **developers** (e.g., **S Molly**) implement the systems and algorithms. In doing so, it is expected that they know not only the technical details but also being aware of ethical implications. In contrast, the **management** has to mediate between the **investors/shareholders**, usually following a profit interest, the cost of implementing ethically flawless systems, and the pressure by the customers (police) to develop usable, efficient and productive systems. It is important to note that the software provider has typically no complete control over all aspects of a software system, as typically external dependencies, models, or training data are being used.

Civil Society — Civil society can materialize not only in the form of mass media through its **members**, but also in the form of NGOs and human rights groups, arts and culture, street protests, whistleblowers, ethics councils, and so on. As such, it deliberates on what can be considered as acceptable ethical behavior in a given society, which parliament follows (through elections), and which can change over time. On the one hand, civil society can act as a corrective, for example, through critical reporting by the **press**, or legal advocacy through **NGOs or human rights groups**. Cases of unfair treatment, when entering the public agenda, can trigger the revisiting of fundamental ethical questions (as in the case of the criticism of the Northpointe recidivism algorithm and the debates it triggered about different notions of fairness and justice [LA16; Spi17; Gre18]). However, mass media and public deliberation can also proliferate misleading ideas about what algorithms can and cannot do. These "socio-technical imaginaries" [Jas15] have concrete implications for how systems are being used, for the transfer and negotiation of responsibility, as well as public acceptance [BK21].

3.4 Mitigation Techniques through Visual Analytics

Addressing the ethical issues raised at the outset of this chapter and negotiating the conflicting interest of different stakeholders is not a trivial task. Given all the different stakes involved, ethically-aware design of intelligence applications can not reasonably aim at implementing technical solutions to safeguard against all possible pitfalls. Rather it seeks to accomplish a serious consideration and balancing of the inherent trade-offs and inter-dependencies between different concerns, interests, and principles. For example, privacy-by-design may limit possibilities for advanced accountability. It thus needs to be negotiated which good is more important in each specific context and how to best achieve this. In doing so, we propose a socio-technical approach, not looking at possible technical solutions in isolation but as embedded phenomena, interacting with their environment [Suc07]. Human interaction with technology is shaped by increasingly sophisticated and environmentally interwoven interfaces, connecting technical components, humans, and their surroundings [PVV20]. Despite a rising appreciation of milieu-oriented approaches to understanding and designing interfaces [AGV19], the development of interfaces has traditionally been dominated by technical disciplines and, from an ethico-political point of view, not sufficiently considered the complex socio-technical dynamics at play [Hoo14; Gal12]. To fill this gap, it seems most productive to work with an extended definition of the interface going beyond isolated, technocentric meanings [LHKP22; Lip17]. We thus adopt the proposal to focus on "interfacing" as a joint practice of "becoming-with" of humans, machines, and environments [Har17]. One could also call this process an intra-action [Bar14], in which something new emerges, irreducible to its parts. Such an understanding leads to a more comprehensive understanding and appreciation of what needs to be considered when designing user interfaces, especially in sensitive high-stakes areas such as communication analysis of intelligence information.

3.4.1 Technical Measures

In the following, we investigate common interface design methods employed in VA-based systems—the technical side—from an interfacing perspective—including the political, responsibility, intra-active, and ethical—as part of an intertwined becoming-with one another. When describing relevant aspects, we refer back to the actors from the scenario in Section 3.3.3. As such, we identify areas where VA, through its human agency approach, is superior to fully-automated systems while also considering the additional burdens through the distribution of responsibility.

Interactive Exploration — One key concept of VA is that instead of merely generating results, such a system supports the knowledge generation process of analysts by enabling them to learn from the data space through supported **interactions** [KAF⁺08]. As part of this process, ethical mitigation techniques can be integrated. In the context of communication analysis, we introduce some of the critical aspects of VA by the example of HYPER-MATRIX [FAS⁺20], a conversation topic analysis and probability framework that uses a geometric deep learning approach. In our hypothetical scenario, police analyst **S Martin** uses the software. Instead of presenting lists of users and topics, it features an interactive design, shown in Figure 3.2.

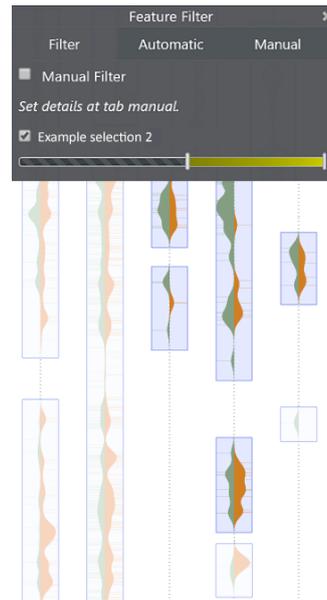
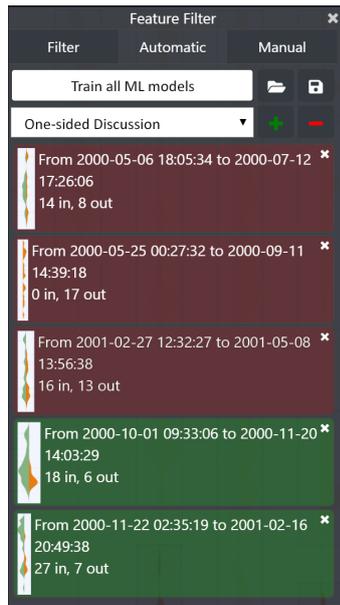
The developers, in our scenario **S Molly**, deployed a **matrix-based visualization** **B** for the hypergraph network structure due to increased scalability, which represents model predictions through **clustering** and **color-encoded** by confidence. The design can be considered as a form of **dimensionality reduction** [KP11], presenting the complex tensor model in a more comprehensible format. This supports the detection of patterns, while color-encoding facilitates pre-attentive understanding, helping **S Martin** to distinguish between users communicating about similar topics, like **S Chris** and **S Carlos**. Further, a multi-level visual **semantic zoom** through multiple, more detailed **in-line visualizations**, shown as insets **D**, allows for a more-detailed exploration, preventing an initial mental overload of **S Martin**. **Steering** is offered by interactive control elements **A**, allowing **S Martin** to set methods, cutoffs, and thresholds, thereby granting him agency and creativity in his usage of the system. Similarly, the system features elements from **active learning** enabling **S Martin** to interactively modify the model **C** to create something new and unique for the purpose at hand in the spirit of his analysis “becoming-with” the system. When using the system, the analyst explores the probabilities, refines model parameters, investigates hypotheses, validates change effects as part of an (indeed intra-active and) iterative analysis **feedback loop**.

The Machine Side - Analysis and Active Learning — An example of an intra-active becoming of investigator and machine is active learning. Figure 3.3 shows how an analyst provides labeled examples to the system improving its probabilistic accuracy. Labeling everything is tedious and time-consuming when done manually by the user. Here, **intelligent labeling** techniques can help by only requesting human input when required, relieving analysts from exploring basic or irrelevant patterns [Cor19]. This concept can be applied **universally** to any number of **feedback mechanisms** between system and user, affecting data selection, machine learning models, heuristic algorithms, or their parameters. Through active learning, **S Martin** can integrate its experience-saturated as well as its domain-specific expert knowledge into the analysis process, thereby sharpening the analysis result with regard to

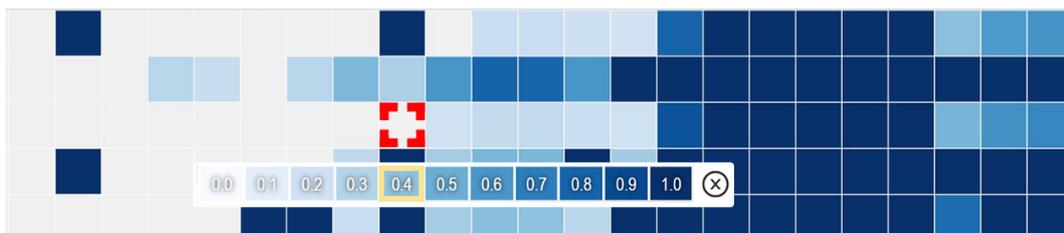
the field's unique requirements. Methods such as color coding imply a form of nudging through preattentive processing. Transparency, in turn, can be conveyed by visualizing consequences and effects of the active learning approach (e.g., which entries are subsequently classified differently and by how much). This corresponds to a "what-if" preview, which supports the selection process, but can also act as a "control", e.g., against unintended side effects.

The Human Side - Guidance and Explainability — Guidance describes the interplay between system and user actions and their understanding in the context of machine learning, explainable artificial intelligence (XAI), and knowledge generation. One form of guidance can manifest through the system: it can, depending on the context, nudge the user in the right direction, for example, by showing similar matches or conflicting options. In the context of learning and teaching, it exhibits a wider dynamic, encompassing **system teaching, user teaching, system learning, and user learning**. As a process, it can be described by the knowledge generation model [KAF⁺08; SSS⁺14] and by the co-adaptive guidance process [SJB⁺20; SJB⁺21]. Different forms of guidance can be achieved by a **visually abstract** visualization as well as **conceptual user interaction** design. For visualization, **abstract representations** like glyph [FIBK17] can be used for improved recognizability and comparability. In communication analysis, commonly used representations are text, highlighting, concept extraction, and network display. However, depending on the individual needs, they may not suffice. During the design, the aim of the **representation**, the selection of the appropriate **visualization technique**, the **visual variables** employed, and the **color-schemes** used has to be considered. **Inherent biases** can play an essential role, affecting values as social biases (e.g., homogeneity bias), actions (e.g., blind spot or Ostrich effect), and **perceptions** (e.g., illusions or Weber-Fechner-Law) [Ell18], both during development as well as usage. For interaction, instead of filtering communication texts by keywords (the selection of which might be biased or incomplete), a **visual query language** [FSS⁺21] can be used where conditions are based on concepts. Here, the system may similarly suggest additional concepts for differentiation or indicate a too restricted search. In implementing such interfaces, care has to be taken regarding a neutral representation while considering the levels of detail and abstraction [JSS⁺18]. Too much abstraction can lead to a loss of context.

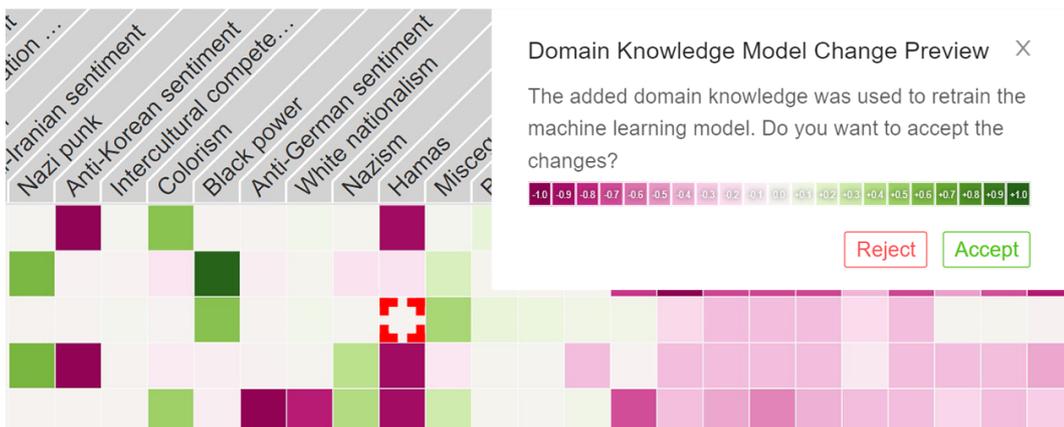
Provenance — As systems become more complex and the number of available interactions increases, the number and sequence of necessary steps during the investigation expands rapidly, making it increasingly difficult to fathom and explain them after the fact. In this context, while many systems focus on supporting the process of knowledge generation, few also emphasize how this process is carried out. [KM13; LAH20]. For reproducibility, which is crucial for accountability (see Section 3.4.2), knowing how analysts like **S Martin** used the system to draw conclusions



(a) Manually supplying positive (green) and negative (red) examples to the system during active learning. (b) Classification model applied to communication episodes after active learning improved filtering.



(c) Manually changing the connectivity strength for a hypergraph model entry through selection, requesting a retraining of the contained machine learning model.



(d) Visualized changes in the model prediction (from negative to no to positive change). Before the changes are either rejected or accepted, the system nudges the analyst to examine the results for unintended side effects.

Figure 3.3: Examples of active learning in Conversational Dynamics [SFS⁺19] (a, b, see Chapter 6) and HYPER-MATRIX [FAS⁺20] (c, d, see Chapter 5).

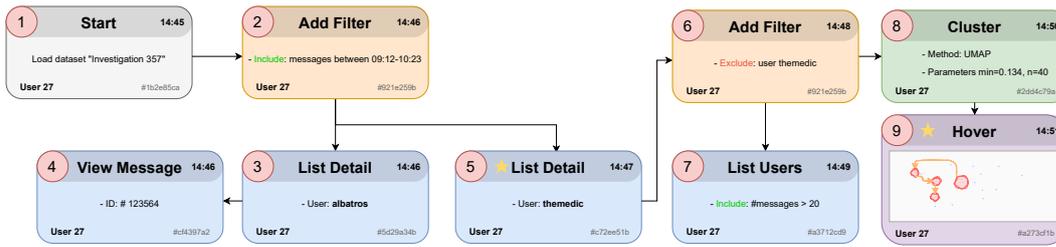


Figure 3.4: Example sketch of a provenance history component using a directed acyclic graph (DAG) approach. Shown are individual branching steps with inline details instead of a linear history, allowing for a more complete picture of the explored steps (1-9) and reverting from dead-ends (3,4,7).

is essential, but this becomes increasingly complicated when iterative and intra-active processes are involved. Instead of a **linear timeline** of steps or a **list** of explored hypotheses, VA applications can store interaction chronicles **in context** with the data, remembering **settings, views, and actions**. Through techniques like time-stamping, hashing, and digitally signing information, the chain of evidence becomes proof-able. As such, they can provide tamper-proof **tracking** as well as **replay** functionality to revisit intermediate steps, while also enabling approaches like **provenance graphs** as shown in Figure 3.4, strengthening the chain of evidence for both analysts and subjects.

3.4.2 Balancing Advantages and Risks Through the Interface

In the following paragraph, we identify advantages (**A1–A6**) but also the limitations and risks involved in using visual analytics for the analysis of intelligence data. To do so, we refer back to the various challenges and conflicts identified in Sections 3.2 and 3.3.

A1. User Agency — One key benefit of Visual Analytics (VA) methodologies is enhancing user agency, raising awareness against unverified, machine-generated results. Preventing this blind obedience is particularly important due to the often unrealistic expectations towards automatic solutions that many users exhibit. VA rectifies the imbalance and levels the playing field between user and system by forming a con-genial joint agency that promotes collaborative growth instead of displacement or rivalry. This can be attributed to the approaches inherent design favoring cooperation over a built-in hierarchy. Moreover, in scenarios where human analysts might deviate from system-suggested outcomes, instead of fostering a binary decision-making process, the concept of **explainability** can enable analysts, like **S Martin**, to reconsider their own actions and inputs that impact the outcomes of the system through active learning. This can eventually contribute to pinpointing dead ends in the analytical process.

A2. Privacy — Guidance mechanisms enable the partial avoidance of privacy concerns as selective presentations permit system-based management, for example, in relation to uninvolved third parties, ensuring the personal life of individuals, like Sarah, remains protected from actual humans if not unmarked through active confirmation.

A3. Fairness — Compared to systems trained exclusively on past criminal records, employing visual analytics for dealing with heterogeneous data has advantages: The former often contain harmful racial, gender, and other biases, thus reinforcing historically developed power structures [LI16]. VA empowers analysts to scrutinize the reasoning process and apply corrective actions using their domain knowledge or common sense.

A4. Efficiency — Big data-oriented VA can illuminate operational dynamics, unveiling previously undetected correlations [MC13]. Beneficial **abstractions** and robust **analytical capabilities** can be suggested by visual analytics systems, efficiently assisting investigators in identifying subtle patterns, the literal needle in the haystack. Nevertheless, each design developed for enhanced efficiency needs to be evaluated for possible risks regarding misinterpretation and oversimplification.

A5. Literacy — VA motivates users to actively and innovatively interact with and utilize the systems, enabling advanced literacy through consistent practice. This acquired literacy can then be disseminated among colleagues and other users, for instance, via the integrated **sharing of recipes** through VA systems (e.g., *Common actions here are...*).

A6. Customization — The development of bespoke solutions over one-size-fits-all approaches is supported by active learning. Generic one-fit solutions often disappoint and under-perform in combating organized crime, necessitating highly localized and specialized strategies, alongside in-depth **domain knowledge** of seasoned, highly skilled experts [Pao02]. As intentionally employing a nonlinear, swarm-like approach is the exact strategic tactic of organized crime, this requirement is no coincidence but rather a typical artifact. They aim to confound investigators and thwart pattern-based strategies, which equally apply to systems. In such scenarios, the input of seasoned experts familiar with these tactics is paramount. Hence, technical solutions that **combine** efficiency-boosting **automation** with **human** criminological **expertise** to create custom solutions hold the potential for significant advancements. Moreover, it is vital for systems to possess the capability to adapt dynamically to evolving conditions and environments.

3.4.3 Risks, Limitations and Additional Measures Required

Similarly, we identified risks (**R1–R5**) and limitations and imperative measures:

R1. Lack of Accountability — Accountability (C6), which virtually always requires reproducibility, presents a substantial challenge. This is because visual analytics uses interfacing, where attributing responsibility for possibly flawed outputs becomes exceedingly complex: Due to the intertwined nature and the joint becoming-with one another of an analyst’s knowledge and the system’s analysis, delineating and attributing cause and effect is nearly impossible. The elements to consider are not merely the precise *software versions*, the initial *seeds* of pseudo-random generators for non-deterministic algorithms, or the exact *data utilized during training*. Detailed information regarding how analysts, such as **S Martin**, **interacted** with the system is equally significant. Merely presenting the final analysis is insufficient, given that **S Martin**’s decision-making process was not strictly sequential but influenced. Learning and refinement occurred throughout the process, and someone in a reviewing role, like judge **S Muller**, would want assurances that all plausible hypotheses were considered rather than neglected or even purposefully ignored. Hence, advanced **provenance** methodologies in VA designs are essential for providing accountability.

R2. Training and Community-Building Among Users — The assumption that users like police officers inherently possess the requisite knowledge to comprehend the particularities of ML and VA systems, their potential limitations, and ongoing development are certainly misplaced. Developers, for example, **S Molly**, face the challenge of expending considerable efforts to provide understandable explanations about the workings of the system while also training users for its effective and critical usage. Ideally, such training should evolve from a unilateral how-to-use a feature to an interactive process involving both software providers and analysts, such as **S Martin**, and also happen *among* the analysts themselves (perhaps through design studies). Potential strategies and directions that can be taken include privacy-preserving **gamification** [SHMK14] elements, for example, practical challenges and tasks, within the **visual analytics system** as well as online discussion **forums** that promote constructive critique and socially negotiate risk and limitations. However, the introduction of new software also alters power structures [Hen91], possibly favoring digitally more literate younger officers and marginalizing veteran investigators [KEL18]. To preserve the invaluable experiential knowledge of these seasoned investigators, it is worth considering technical solutions that incorporate, rather than exclude, these qualitative dimensions.

R3. Prevent Automated Inequality — Active learning may inadvertently enable the integration of individual or institutional racism due to unconscious biases of those influencing and steering the learning process [Ell18]. This risk is particularly critical in collaborative solutions, which transfer substantial training and design responsibility to investigators and officers. Given alarming research on unconscious racial and sexist biases prevalent among police forces [KES20; FBBP16; Fra05; PP05],

this risk requires further counter- and control measures and awareness regarding these potential issues must be raised. **Human in command** solutions involving active learning or automation based on user input, exemplified by **S Martin**, should only be implemented and deployed following extensive anti-discrimination training and education. Similarly, system developers, such as **S Molly**, may also inadvertently introduce biases into the system through a flawed design or through incorrect assumptions. Therefore, incorporating mitigation measures at the VA system level itself is especially beneficial, as it pays off through the disseminator effect.

R4. Facilitating Critical Reflection — The user interface, controlled through the system, should actively encourage moments of reflection [Bau15]. Such reflective moments (e.g., through nudges) can aid users in identifying their unconscious biases during their interaction with the system. For example, warning signs can be displayed after a user repeatedly performs allegedly racist or sexist search queries before such queries are blocked behind confirmation prompts [WSE19] or requiring secondary sign-offs. These prompts, of course, have to be accompanied by transparent explanations. On the opposite, reflective moments may also enable users to challenge the system's outputs and decisions [WSE19]. Here, investigators support a critical analysis of the software itself. By incorporating **reflective design** [SBDK05], users are better equipped to both interpret the active learning process and evaluate system results, thus fostering a critical use of the software as well as improving technical literacy.

R5. Human Oversight — In light of the risks discussed above, support for shared responsibility and joint agency between analysts and systems is crucial but ultimately must be supplemented with additional layers of human oversight. Regular **reviews** of system training are essential for ensuring fairness, especially towards protected (minority) groups. An effective VA **provenance** system can track individual users who interact and incorporate changes through active learning while simultaneously safeguarding their identities. This facilitates a critical discussion of problematic decisions among colleagues and authorities in a safe environment without the fear of repercussions. Leadership roles, such as **S Dr. D**, could then validate the extent of bias within their organizations in a privacy-respecting manner. However, care must be taken not to create a culture of control where experienced officers feel under constant surveillance on their own and stripped of their agency [KEL18]

3.4.4 Negotiating Risks and Advantages Through the Interface

Visual Analytic systems can be considered a distinctive type of interface-oriented solutions, representative of a data infrastructure platform that facilitates the acquisition, manipulation, utilization, and storing of information, along with associated metadata, sourced from heterogeneous origins. Such interactive, user-centered real-

time approaches have the potential to counterbalance inherent biases concealed within fully automated systems due to being trained on imbalanced, problematic, or "dirty" data or datasets that are otherwise insufficient. Yet, inadvertently, these systems can also accentuate discriminatory biases, thereby increasing the responsibility of the users. As a consequence, the implementation of in-process fairness measures has become increasingly necessary to counteract these issues. Similarly, ensuring traceability and untampered provenance of the genesis process is extremely relevant. This enables accountability in instances of misuse and, when fairness violations are detected, to trigger appropriate retraining procedures. Further, the demand for transparency must be balanced with privacy considerations, assuming a certain degree of abstraction. Nevertheless, operational success still relies on adequately trained personnel, at least to some degree. Compared to manual analysis through separate modalities or a fully-automated process, interactive visual analytics can be designed with built-in ethical considerations, encompassing the entire knowledge generation process. This approach considers and manages user expectations in line with technical capabilities (C4), counters system opacity (C3), and adheres to privacy requirements (C2). Multiple stakeholders exercising human oversight (C5) can help identify instances of discrimination (C1). Concurrently, the automatic collection of tamper-proof provenance about both the system and the human operators enables visual analytics systems to bolster trustworthiness and accountability (C6).

3.5 Conclusion

We have shown in detail how various VA methods can address ethical challenges in advanced analytical systems. While we have highlighted concrete points to consider, we do not intend to present a fixed set of rules for the ethically conscious design of VA systems. Indeed, in our view, this is always a matter of negotiating trade-offs between conflicting interests, which can vary widely depending on the unfolding interaction dynamics. Therefore, we aim to stimulate a discussion about the consideration of ethical implications as an integral part of the design process from the outset.

Although we focus on the case of intelligence applications, many of the results of our work and the ethical discussion are more generally applicable to the design of VA applications. This is because, on a more abstract level, our approach leads us to the insight that one of the main advantages of VA methods is that they take the *interface* as a starting point for technological innovation. This means that innovation is approached not only in a technical, but rather in a **socio-technical** way. Useful innovations cannot exist in isolation and should pay attention to their impact on communities and society as a whole. This shifts the focus to embedding technologies

into existing social institutions such as police departments and criminal justice agencies. Here, new technologies are most useful when combining the benefits of efficiency-enhancing automation with the experience-saturated knowledge of investigators. Concentrating on the interface has obvious advantages, as it implies a focus on the situated nature of human-computer-configurations.

By capitalizing on instead of trying to stabilize the potential openness, and thus eventfulness of technological interconnectedness, undesirable dynamics can be dealt with much more proactively, while flexibly addressing the situational and individual needs of different cases. The issue is not only what interacts with whom, but also how new phenomena emerge as part of complex intra-active configurations of people and automation systems. Discriminating behavior of individual investigators, for example, might not affect the functioning of the police force as a whole. When multiplied and cemented in automation processes, however, it can contribute to structural discrimination and patterns of unfair treatment on a larger scale. By responding transparently to investigator input, VA can help well-meaning analysts recognize their often unconscious biases by showing them how harmful social stereotypes sometimes cause them to overlook features and stumble down blind alleys.

To avoid risks and increase benefits, it is important to keep in mind that interfacing is not simply the matching of two separate entities, but the creation of something fundamentally new, a hybrid of human and machine. This hybrid requires tailored quality insurance measures such as adapted training, new forms of oversight involving technical, legal, and ethical experts, and also adapted policy and ethical frameworks that focus not solely on the technologies or on the user, but on what emerges as something new in the interaction between these entities.

Part II | Identification and Interpretation

If it disagrees with experiment, it's wrong. In that simple statement is the key to science. It doesn't make any difference how beautiful your guess is, it doesn't matter how smart you are, who made the guess, or what his name is. If it disagrees with experiment, it's wrong. That's all there is to it.

— **Richard P. Feynman**, *Physicist*

4

Communication Pattern Identification and Interpretation: A Survey of Hypergraph Model Visualizations

When investigating the communication behavior of more than two parties, the participants can be modeled as a (social) network. Such a network is usually represented through the use of graphs. However, regular graphs encounter difficulties accurately reflecting multi-party communications about multiple topics in a complex context, which is the default case in any reasonable real-world scenario. These shortcomings of graph modeling are common for polyadic interaction processes in a wide range of domains, ranging from gene interactions in biology to traffic networks. As such, these processes can be modeled more precisely as (temporal) hypergraphs than as regular graphs. This is because hypergraphs generalize graphs by extending edges to connect any number of vertices, allowing complex relationships to be described more accurately and allow for a more detailed prediction of their behavior over time.

Contents

4.1	Hypergraph Visualizations	90
4.2	Related Work	91
4.3	Methodology	92
4.4	Literature Survey	94
4.5	Comparison and Analysis	98
4.6	Discussion and Future Work	100
4.7	Conclusion	101

However, interactive visualization methods for hypergraphs and hypergraph-based models have rarely been explored or systematically analyzed. This chapter reviews the existing research landscape for hypergraph and hypergraph model visualizations and categorically analyzes the existing approaches. With respect to comparison criteria, we focus on application scenario—both conceptually and from a technical side—scalability, interaction support, and evaluation methods.

This chapter is based on the publication [FFKS21] and major parts of the following sections have appeared in:

- [FFKS21]: **Maximilian T. Fischer**, Alexander Frings, Daniel A. Keim, and Daniel Seebacher. “Towards a Survey on Static and Dynamic Hypergraph Visualizations”. In: *2021 IEEE Visualization Conference (VIS)*. IEEE, 2021, pp. 81–85. DOI: [10.1109/VIS49827.2021.9623305](https://doi.org/10.1109/VIS49827.2021.9623305).

For a statement of the scientific contributions, as well as the division of responsibilities and work in this publication, please refer to Chapters 1.2 (p. 22ff) and 1.3 (p. 26ff), respectively.

4.1 Hypergraph Visualizations

Hypergraphs are an extension of graphs, allowing edges to connect more than two nodes. Mathematically, a graph $G = (V, E)$ is defined as a pair with n vertices $V = v_1, \dots, v_n$ and a set of paired vertices $E \subseteq \{(a, b) \mid (a, b) \in V^2, a \neq b\}$ forming the graph edges. In contrast, an undirected hypergraph $H = (V, E)$ is a pair where the vertices (defined equally) are connected through a multi-set $E = e_1, \dots, e_m$ that form m distinct hyperedges [Ber84], which may connect arbitrary vertices (usually ≥ 2 vertices). This is a desired behavior for modeling complex, multi-entity processes as it allows to capture mutual dependencies more accurately and logically separated [KHT09], efficiently encapsulating group dynamics [VBP⁺19]. Especially for group structure capturing and sub-group combinations, as they can occur in the context of human communication and social media data, hypergraphs can provide advantages [OSH⁺07]. As such, the modeling through hypergraphs has received increasing research interest in the last few years [WXLW07; VBW17], as continuing with the general concept of (hyper-) graph modeling allows to continue using concepts like edge attributes or weights and many graph algorithms.

However, the more complex representational difficulty of hypergraphs brings along several visualization challenges [HC14a]. Traditionally, Venn or Euler diagrams

have often been employed to visualize hypergraphs [JP87; Mäk90; BE01]. However, these conventional visualization methods often use color to differentiate between hyperedges, which limits their scalability significantly. Consequently complex or large-sized hypergraphs are virtually impossible to meaningfully visualize using Venn or Euler diagrams.

Even with these challenges existing, recent developments have shown an increased interest regarding network relations and their temporal development, giving rise to the concept of temporal (or dynamic) hypergraphs. This extension further adds a temporal dimension to (hyper-) graphs, allowing for the comparison of network states at different time intervals. As such, they pose an additional challenge for the employed visualization techniques to capture both the structural intrinsics as well as the temporal dynamics.

Despite the overall advancement in this research area, elaborated hypergraph visualizations are still relatively novel and not thoroughly explored yet. There is a lack of detailed comparisons between different approaches, and this gap underscores the need for a systematic exploration of hypergraph visualization methodologies to be aware of their respective strengths and weaknesses. To address these issues, we present a survey of hypergraph visualizations, making the following contributions:

Contributions

- A systematic **literature review** of existing approaches for static and dynamic hypergraph (model) visualization.
- A methodology for **comparison criteria** between hypergraphs, critically assessing the different approaches.

With this survey, we aim to provide insights into the existing research landscape of generic hypergraph visualizations.

4.2 Related Work

Hypergraph visualizations, despite being known for some time in the form of Venn or Euler diagrams, have so far not been systematically explored and lack a comprehensive review of the methodologies employed, alongside their respective advantages and limitations. Due to its partial relatedness to set- and association graph visualizations, we also consider such surveys as a starting point. Relevant works have been produced by Alsallakh et al. [AMA⁺16] and Chen et al. [CGZ⁺19] for sets, while several overview reports exist for graph visualizations [VBW17; BBDW17].

Alsallakh et al. [AMA⁺16] provided a seminal paper on visualization options for sets. Their work covers a variety of approaches, but many are adaptations

or enhancements of Euler and Venn diagrams, such as BubbleSets [CPC09]. It is important to note, however, that not all of these methods are applicable to generic hypergraphs due to their domain-specific nature. Chen et al.'s survey [CGZ⁺19] instead explores association relationships in graphs. Nevertheless, it still retains a focus on set-relations. They provide a balanced analysis of different visualization technique and focus on relational data. As part of their survey, they investigate two visualization methods for hypergraphs: first a radial layout designed by Kerren et al. [KJ13], and second a fixed-node visualization presented by Xia et al. [XLH⁺11].

In relation to graph visualizations, the survey by Vehlow et al. [VBW17] presents a visualization technique taxonomy. They differentiate between visual node attributes, superimposed, juxtaposed, as well as embedded visual styles. The shift towards dynamic graphs highlighted before becomes more noticeable in recent surveys, such as the one conducted by Beck et al. [BBDW17]. They categorize dynamic graph visualizations into animated node-link diagrams and timeline structures, noting a trend towards the use of timelines in the literature.

While some concepts and taxonomies can be applied to hypergraphs, shortcomings of these existing surveys relate to their relative old age, their incompleteness and their lack of addressing more recent hybrid approaches [SAKW19; VBP⁺19; FAS⁺20], which do not always fit well within the established criteria. These more recent approaches, like PAOHvis by Valdivia et al. [VBP⁺19], the first representation of dynamic hypergraphs, a hybrid visualization for dynamic hypergraphs by Streeb et al. [SAKW19], as well as HYPER-MATRIX by Fischer et al. [FAS⁺20] underline the importance and potential of further research in this area.

4.3 Methodology

To analyze state-of-the-art approaches on hypergraph visualizations, we conducted a search of relevant literature before collecting further approaches via cross-references. We aim to analyze visualization literature that applies to hypergraphs as well as hypergraph-based models, and those techniques that focus on representing the inner connectivity of hypergraphs, which can be leveraged in the visualization.

4.3.1 Source and Selection Methodology

Orienting ourselves on semi-automated literature surveys, such as the one of Sacha et al. [SZS⁺17], we started the paper collection process using a keyword-based search for “hypergraph”, as well as common variations like “hyper-graph, hyper graph, ...”. To further streamline the focus of our survey, we consider only approaches from the following high-quality journals and conferences within the last 23 years (since 2000):

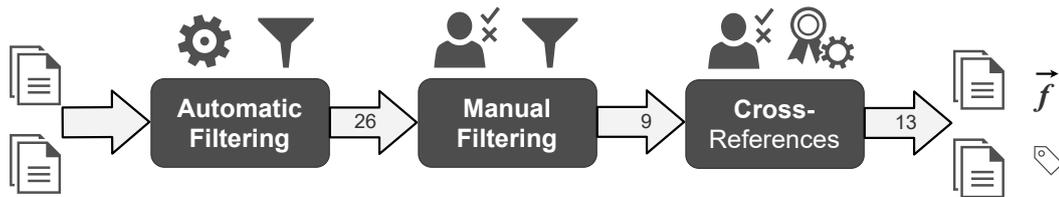


Figure 4.1: The paper collection process consists of three main steps: (1) Automated filtering, (2) manual filtering, and (3) cross-referencing.

- IEEE Transactions of Visualization and Computer Graphics (**TVCG**, including the **IEEE VIS** proceedings)
- Computer Graphics Forum (CGF, including the **EuroVis** proceedings and EuroVA)

For the actual paper selection methodology, we follow a three-step approach (see also Figure 4.1). Automated filtering resulted in five (TVCG) and 21 (CGF) approaches; manual filtering reduced this first to three and ten, while in a second iteration after careful considerations, to three and six, i.e., nine approaches. Using cross-references and survey data, we identified an additional four relevant approaches, leading to 14 approaches in total, which we compare using the following comparison framework.

4.3.2 Comparison Framework

We describe seven comparison criteria covering some aspects common to graphs and some favoring aspects specifically designed for hypergraphs. We curated this selection based on an adapted combination of criteria based on existing surveys [VBW17; BBDW17] as well as aspects discussed in the literature [VBP⁺19; FAS⁺20], and complemented them by additional distinguishing criteria.

Representation Method The representation method determines the basic visualization type, where existing approaches can be classed as either *node-link* diagrams, *timeline-based* techniques, or *matrix-based* approaches.

Scalability Scalability varies significantly between the different visualization approaches. For easier comparison, we chose five distinguishable scalability levels: very low scalability (■□□□, less than ten nodes or hyperedges), low scalability (■■□□, between 10 and 50), medium scalability (■■■■□, between 50 to 200), high scalability (■■■■■, between 200 and 1000), and very high scalability (■■■■■■, more than 1000). Many basic representations are only applicable when there are only a limited number of vertices and hyperedges. Further, they are prone to clutter for a higher count of edges, whereas different techniques perform better.

Static vs. Dynamic Any meaningful analysis of dynamic processes usually requires support for displaying dynamic (i.e., temporal) data (■) in contrast to static (□) information. The methods employed in the presented approaches differ widely,

but many are still limited to some degree. Options used are highlighting employed at different stages, direct comparison, but also just a simple timeline.

Interactivity Some visualizations only represent a static rendering and do not offer any kind of interactivity (□), while others allow the visual representation or the underlying model to be tuned (■) and adapted interactively, for example through filtering, rearranging, or highlighting. This can be leveraged for a coherent representation, guide the analysis workflow, or represent information in a comprehensive framework.

Tasks The strength of an approach can often be better judged by a description of the tasks that it aims to support or the presentation of possible use cases in a specific domain. If the corresponding paper contains such tasks definitions or describes possible usage scenarios, we note this as part of this category.

Evaluation Evaluating the design concepts together with experts has become increasingly important in recent years. This avoids the creation of problematic approaches that encounter usability and applicability issues in practice. Such evaluations can either come in the form of comparative evaluation (↔) between related approaches, a convincing case study describing a possible domain-specific application scenario in detail, as well as an user-based qualitative or quantitative evaluation study. Conducting no evaluation at all (-), which was common a few years ago, is relatively rare by now. Some papers describe case studies (E), while others conduct user studies (P) with participants. In many cases the users interviewed are domain experts and the studies are often qualitative evaluations. There, they conduct analytical tasks, while quantitative evaluations are rare.

4.4 Literature Survey

In the following comparison, we group the approaches by their representation method, while the order within each category is chronological.

4.4.1 Node-Link-based Approaches

In the following, we describe node-link-based approaches, which are visually depicted in Figure 4.2.

Bubble Sets Collins et al. [CPC09] use isocontours which are overlaid over connected existing elements (see Figure 4.2a). Conceptually, this is very similar to the colored hulls used in an Euler diagram. However, to generate the isocontours, they employ a marching square algorithm. The visualization itself is static, but nodes can be moved and added or deleted.

Software Artifact Hypergraph Visualization Kapec et al. [Kap10] describe a software development and programming environment that represents individual

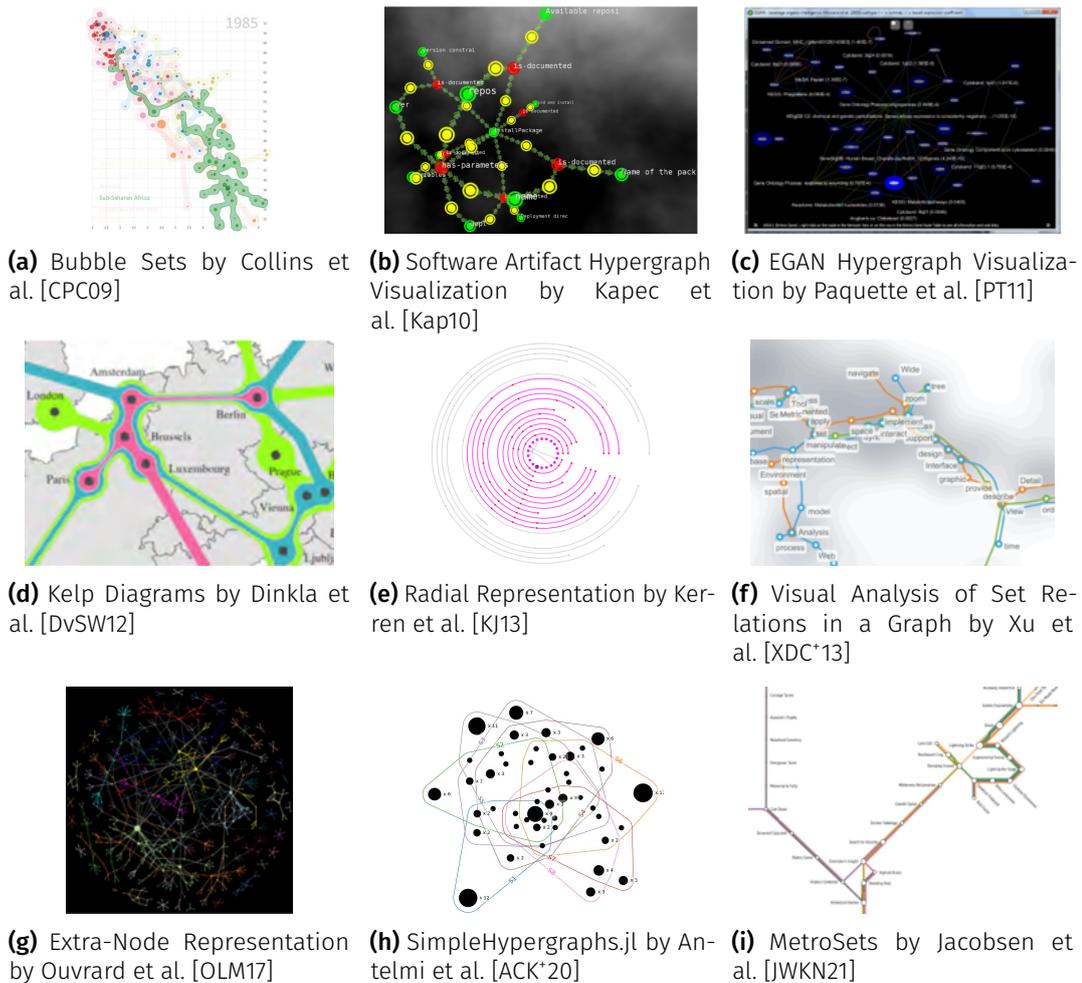


Figure 4.2: Node-link-based approaches for (temporal) hypergraph visualizations.

functions, their source code. To enable efficient navigation and understanding of the relationships, the software artifacts mutual dependencies are modeled through hypergraphs (see Figure 4.2b). These are represented using a force-directed 3D layout with colored spheres to differentiate between callees and callers.

EGAN Hypergraph Visualization Paquette et al. [PT11] present an extensible visualization for sorting and classifying gene lists (see Figure 4.2c). The method which they named Exploratory Gene Association Networks (EGAN) is based on a node-link-based approach but introduces a distinct association meta-node that connects with each relevant entity, similar to the concept of hyperedges. Consequently, this method can enhance the scalability in comparison with color-based methods, while still maintaining similarities to node-link representations.

Kelp Diagram Dinkla et al. [DvSW12] augment the idea of colored-hulls and bubble-maps through the utilization of overlapping lines, node coloration, and vari-

able link thicknesses (see Figure 4.2d). Their approach is primarily non-interactive and is intended for stationary geospatial applications due to its constrained scalability. An evolutionary successor, KelpFusion [MRS⁺13], aims to mitigate clutter through a combined hull and linear set representations as well as enhanced computational efficiency.

Radial Representation Kerren et al. [KJ13] introduce a novel layout paradigm in which hyperedges are depicted as circular dotted lines encircling centrally positioned nodes (see Figure 4.2e). This arrangement eliminates overlaps and can enhance the scalability, but still represented a constrained layout of node-link connections. The approach offers several interactive features, including filtering, highlighting groupings, and link modification. Some aspects like labeling or graph comparison are missing, but discussed as potential extension. The approach itself was evaluated through a small-scale user study.

Visual Analysis of Set Relations in a Graph Xu et al. [XDC⁺13] base their visualization on classical node-link diagrams, which they extend through a glyph-based overlay technique (see Figure 4.2f), while also using colored connectors. While they aim to visualize the set relations and node (path) distance in graphs, no interaction concepts are present and scalability of their approach is severely limited.

Extra-Node Representation Ouvrard et al. [OLM17] introduce an enhanced approach for converting hypergraphs into node-link representations by introducing artificial so-called extra-nodes (see Figure 4.2g). These extra-nodes consolidate and merge multiple edge connections, thereby preserving increased context while simultaneously reducing visual clutter. The applicability of their approach is shown through an extensive user study.

SimpleHypergraphs.jl Antelmi et al. [ACK⁺20] developed SimpleHypergraph.jl as a hypergraph visualization framework (see Figure 4.2h). It uses a combination of Julia, Python and D3.js, supports only limited interactivity, and the overall idea is slightly reminiscent of EGAN [PT11]. From a visualization perspective, a convex hull encompasses nodes, while hyperedges are represented as sub-graphs. The main contributions here are in providing a library for drawing hypergraphs with some slightly improved visual properties, but not a completely new visualization design.

MetroSets MetroSets by Jacobsen et al. [JWKN21] focuses on geospatial and train network data (see Figure 4.2i). It is one of the few approaches capable to support temporal data, which is mapped to a time axis. However, the design leads to much visual clutter through the enforced constraint on angular separation. The approach support different interaction modalities like filtering and highlighting, but is missing aspects live free movement.

4.4.2 Timeline-based Approach

In the following, we describe timeline-based approaches, which are visually depicted in Figure 4.3.



(a) TimeSets by Nguyen et al. [NXWW16] (b) PAOHvis by Valdivia et al. [VBP+19] (c) Set Streams by Agarwal et al. [AB20]

Figure 4.3: Timeline-based approaches for (temporal) hypergraph visualizations.

TimeSets Nguyen et al. [NXWW16] introduce an approach which in parts reminds of KelpFusion, using color for hyperedge representation, but involves a stacked timeline view (see Figure 4.3a) with various interaction techniques. A shortcoming is the limited scalability of the approach. The technique itself is evaluated through a case study and evaluated through several users.

PAOHvis Valdivia et al. [VBP+19] present PAOHvis, which employs an ordered timeline (see Figure 4.3b). This timeline follows a grided pattern extending along in vertical direction to represent hyperedges: Hyperedges are therefore visualized through vertical lines connecting multiple nodes. In case the available space is limited, they use drip nodes as placeholders. Their approach is built upon earlier research [VBP+18], uses space rather efficiently, features multiple interaction modalities, and facilitates the analysis of related hyperedges.

Set Streams Agarwal et al. [AB20] describe an approach that uses branching and merging information streams in a timeline view (see Figure 4.3c). This process aims to depict dynamic set membership and is conceptually similar to some braiding statistic visualizations in theoretical physics. The technique allows for different set operations and is relatively scalable. The authors validate this approach through a case study and an expert study.

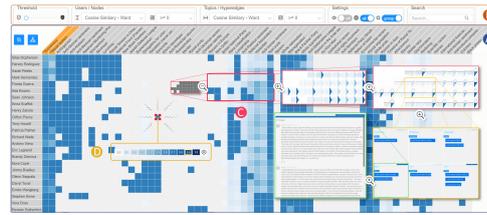
4.4.3 Matrix-based Approach

In the following, we describe matrix-based approaches, which are visually depicted in Figure 4.4.

Visual Analytic Framework for Dynamic Hypergraphs developed by Streeb et al. [SAKW19] present a glyph-based matrix view (see Figure 4.4a) to support temporal data through the use of inline timelines. The scalability is limited and interaction support is basic.



(a) Streeb et al. [SAKW19]



(b) Hyper-Matrix by Fischer et al. [FAS⁺20]

Figure 4.4: Matrix-based approaches for (temporal) hypergraph visualizations.

Hyper-Matrix Fischer et al. [FAS⁺20] propose HYPER-MATRIX (for a detailed discussion, see the next Chapter 5) as the most recent visualization method. It employs a six-level matrix-based representation to visualize dynamic hypergraph structures (see Figure 4.4b). The system internally uses a geometric deep learning model and the visualization supports multiple filtering, grouping, and interaction mechanisms like matrix-reordering, or hierarchical grouping. The approach is evaluated in a comparative evaluation as well as an expert evaluation and features the direct integration of domain knowledge into the underlying machine learning model.

4.5 Comparison and Analysis

In the previous section, we analyzed the current state of hypergraphs and hypergraph models visualizations, which we briefly summarized in Table 4.1. There, we illuminate the shared commonalities between the techniques, while we aim to work out the differences between them in more detail in the following. We observe a clear trend in hypergraph visualization techniques: more recent approaches highlight advanced interactivity and dynamic support. Similarly, conducting user studies has become more prominent over time, reflecting an increasing emphasis on user-centered design and user requirements.

Most hypergraph visualizations are based upon node-link diagrams and employ standard graph drawing techniques through the use of (colored) hulls or with specially crafted nodes. The node-link-based approaches generally can only scale to hyperedge counts ranging from a few (like in Venn diagrams) to several dozens. On the contrary, timeline- and matrix-based approaches demonstrate an increased capacity, managing medium- to large-sized data-sets with several hundred nodes and hyperedges.

A three-dimensional, scalable model is presented by Kapec et al. [Kap10] stands out as an early approach in this direction, which can add some valuable information. However, from the visualization perspective, the different components use colors and shapes overlaid on their sub-graphs to denote node affiliations, which is still similar to the original ideas.

Table 4.1: Overview and comparison of hypergraph visualization techniques, grouped according to primary visualization technique.

	Node-Link			Timeline			Matrix							
	BubbleSets [CPC09]	Kapec [Kap10]	Paquette [PT11]	Kelp D. [DvSW12]	Kerren [KJ13]	Xu [XDC ⁺ 13]	Ouvrard [OLM17]	Antelmi [ACK ⁺ 20]	MetroSets [JWKN21]	TimeSets [NXWW16]	PAOHvis [VBP ⁺ 19]	Set Streams [AB20]	Streeb [SAKW19]	Hyper-Matrix [†] [FAS ⁺ 20]
Scalability	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■	■■■■■
Domain	Geo/Netw.	Softw.	Bio.	Geo/Netw.	Generic	Netw.	Netw.	Netw.	Geo	Events	History	Netw.	Comms	Comms
Dynamic Support	□	□	□	□	■	□	□	□	■	□	■	■	■	■
Interactivity	□	■	■	□	■	□	□	■	□	■	■	■	■	■
Tasks	■	■	■	■	■	■	□	■	■	■	■	■	■	■
Evaluation	ⓔ	-	ⓔ	↔	ⓔ	ⓔ	ⓔ	ⓔ	ⓔ	ⓔ	ⓔ	ⓔ	ⓔ	ⓔ

Legend of symbols used in the table:

- Very Low
- Low
- Medium
- High
- Very High
- Missing
- Partly
- Present

- ⓔ Case Study
- ↔ Comparison
- ⓔ User

[†] HYPER-MATRIX is described in detail in Chapter 5

Node-link methods often come with limitations: Some struggle to present a clear structural overview due to their dependence (overlapping) projections, while others conduct irreversible transformations that can incur an information-loss. Interactive ways to alter the visualization primarily involves changing either the position or the viewpoint of the model. A critical constraint for many methods is scalability, as they quickly become unreadable with increasing complexity. While some approaches try to partly mitigate this through use of artificial nodes or constraint on the arrangement, none can fundamentally escape these constraints, except when involving very particular domain adaptations like Kelp-diagrams [DvSW12].

All advanced techniques include standard interactions such as highlighting. MetroSets [JWKN21] stands out by offering labels for both nodes and hyperedges and providing support for dynamic data.

The most common tasks are the search for nodes in the visualization [Kap10; PT11; DvSW12; KJ13; XDC⁺13; NXWW16; SAKW19; VBP⁺19; FAS⁺20; JWKN21] and the determination of communities/memberships as well as connected components [PT11; DvSW12; XDC⁺13; NXWW16; SAKW19; VBP⁺19; ACK⁺20; AB20; FAS⁺20; JWKN21], sometimes to determine their size [NXWW16]. Other common techniques are membership queries [CPC09; PT11; XDC⁺13; NXWW16; SAKW19; VBP⁺19; FAS⁺20] and the tracing of connections [CPC09; KJ13; NXWW16; VBP⁺19; AB20; FAS⁺20; JWKN21]. This path following is sometimes extended to determine implicit shares and overlap [XDC⁺13] as well as random walks [ACK⁺20]. Less frequent tasks are filtering based on external factors [KJ13; AB20; FAS⁺20], evaluating the readability [NXWW16], classification [AB20], and change detection in models [FAS⁺20].

Matrix-based visualizations are regarded more capable with respect to scalability, which is beneficial for highly correlated datasets. Due to their structured layout, interaction is often more consistent. Among the matrix-based techniques, the PAOHvis approach [VBP⁺19] stands out with its vertical timeline representation of hyperedges, where nodes act as coordinates. The underlying data is visualized through different time-representing layers or by the use of tags. Meanwhile, the most recent technique, HYPER-MATRIX [FAS⁺20], maintains ease-of-use while placing greater emphasis on analyzing connections and correlations. Furthermore, it broadens the concept, providing a generic blueprint for visualizing hypergraph models via an innovative architecture and multi-level visualization technique. This is described in detail in the next [Chapter 5](#).

4.6 Discussion and Future Work

Despite the recent advances in the field, there are still several unexplored opportunities for future research. First, current node-link-based approaches, and even advanced timeline [AB20] and matrix [FAS⁺20] representations, face limitations in

scalability, only supporting up to approximately a thousand nodes or edges. Innovative solutions to enhance scalability, such as the use of extra-nodes [OLM17], the application of aggregation and subsetting techniques, or the development of dense, domain-specific representations, require further exploration.

Second, only a few methods provide support for temporal (dynamic) hypergraphs, and this support is often tailored to specific use cases. In particular, support for many time steps is particularly scarce. Researchers could be inspired by concepts developed for dynamic network visualizations [CSJ⁺20] and explore how they can be adapted and applied in the context of hypergraph visualizations.

Finally, a significant gap is the absence of a standard benchmark dataset and established performance metrics for hypergraph visualizations. This hinders comparison and evaluation of different approaches. Similarly, there is limited discussion [VBP⁺19; FAS⁺20] regarding specific tasks in hypergraph visualizations. Future research should aim to develop standardized metrics, benchmark datasets, and expand the formalization of hypergraph-specific tasks.

4.7 Conclusion

In the last decade, problem modeling through hypergraphs has gained much attention. While some visualization methods have been proposed, there is a noticeable and increasing gap between applied hypergraph research and their visualization. By surveying the existing approaches for hypergraph (model) visualizations, we aim to structure the research space. We first argue for the relevance of such a survey by analyzing existing literature before defining a reproducible paper selection process. Then, we systematically structure comparison criteria before presenting the techniques individually. We discuss the particularities of each technique individually, before coding it, and finally discuss in detail the observations in relation to the other approaches. We find that many visualizations do not leverage the full potential of hypergraphs and are limited in scalability, interactivity, or the support of dynamic hypergraphs. Of the three most promising and generic techniques, PAOHvis [VBP⁺19] and Set Streams [AB20] are both timeline based. However, matrix-based approaches offer unparalleled opportunities regarding scalability, which we address as part of HYPER-MATRIX [FAS⁺20], discussed in detail in the next [Chapter 5](#).

By filling this gap with an overview of hypergraph visualization methods, we aim to provide researchers with a standard reference, promote areas for future work, and set the baseline for a more in-depth survey on hypergraph visualizations.

The great enemy of communication, then, is the illusion of it. We have talked enough: but we have not listened. And by not listening we have failed to concede the immense complexity of our society—and thus the great gaps between ourselves and those with whom we seek understanding.

— William H. Whyte, Sociologist

5

Communication Pattern Identification using Hyper-Matrix

Leveraging hypergraph structures to model advanced processes has become increasingly popular in many domains over the last few years. Regular graphs can encounter difficulties in modeling multi-party communications about different, partly overlapping topics in a complex context. Hypergraphs, however, can provide more precise modeling and representation of the underlying processes while simultaneously reducing the total number of edges needed for describing polyadic processes. As we described in the previous chapter, there is a lack of visualization approaches for hypergraphs that scale beyond a few dozen entries, and matrix-based visualizations have been ill-explored in this regard while offering promising opportunities. Further, the interactive exploration and seamless refinement of such hypergraph-based prediction models still pose a major challenge. In this chapter, we describe HYPER-MATRIX, a novel visual analytics technique that addresses this challenge through a tight coupling between machine learning and interactive visualizations. In particular, the technique incorporates a geometric deep learning model as a blueprint for problem-specific models while integrating visualizations for graph-based and category-based data with a novel combination of interactions for an effective user-driven exploration of hypergraph models. To eliminate demanding context switches and ensure scalability, our matrix-based visualization provides

Contents

5.1	Introduction to Temporal Hypergraphs	104
5.2	Related Work	107
5.3	Extension of Machine Learning to Hypergraphs	111
5.4	Interactive Hypergraph Model Exploration . .	114
5.5	Case Study: Forum Communication Data . . .	120
5.6	Formative Evaluation	123
5.7	Discussion and Future Work	127
5.8	Conclusion	128

drill-down capabilities across multiple levels of semantic zoom, from an overview of model predictions down to the content. We facilitate a focused analysis of relevant connections and groups based on interactive user-steering for filtering and search tasks, a dynamically modifiable partition hierarchy, various matrix reordering techniques, and interactive model feedback. We evaluate our technique in a case study and through formative evaluation with law enforcement experts using real-world internet forum communication data. The results show that our approach surpasses existing solutions in terms of scalability and applicability, enables the incorporation of domain knowledge, and allows for fast search-space traversal. With the proposed technique, we pave the way for the visual analytics of temporal hypergraph models in a wide variety of domains. In the context of this dissertation, HYPER-MATRIX is primarily aimed at identifying patterns and communication participants as well as topics.

This chapter is based on the publication [FAS⁺20] and major parts of the following sections have appeared in:

- [FAS⁺20]: **Maximilian T. Fischer**, Devanshu Arya, Dirk Streeb, Daniel Seebacher, Daniel A. Keim, and Marcel Worring. “Visual Analytics for Temporal Hypergraph Model Exploration”. In: *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2020), pp. 550–560. DOI: [10.1109/TVCG.2020.3030408](https://doi.org/10.1109/TVCG.2020.3030408).

For a statement of the scientific contributions, as well as the division of responsibilities and work in this publication, please refer to Chapters 1.2 (p. 22ff) and 1.3 (p. 26ff), respectively.

5.1 Introduction to Temporal Hypergraphs

A significant volume of real-world data consists of entities and their relationships and can accordingly be modeled mathematically using graph-based approaches. Such approaches are widely applied in many domains, ranging from natural and social sciences to engineering and business. Examples include modeling biological and chemical processes like protein-protein interactions [Prž11], path-signaling [RTK⁺14] or medical feature selection [SRKS16], relationships in computer [WXLW07] as well as human communication networks [OSH⁺07], or knowledge network exploration in business processes [HB05]. Whereas static graphs can represent the fixed relationships between entities, using an undirected or directed graph as a model, many of

the examples presented above are more accurately described as processes with complex interrelations that may change or evolve. Hypergraph modeling shows superior performance in several neural network classification tasks [FYZ⁺19; JWF⁺19] by capturing interrelated concepts through more specialized edges connecting related concepts semantically. Further, geometric deep learning methods together with interactive visualization can help to more accurately model, predict, and explore the model evolution. Considering, for example, conversations, a topic is a time-dependent grouping encompassing users, which cannot be described using a static graph. This evolution of relations should be modeled by dynamic networks. Compared to regular graphs, using edges or separate node types, such modeling often reflects the actual process more accurately.

Dynamic networks are, however, more challenging to model and have traditionally been modeled as regular, undirected graphs, mainly due to computational and visualization limitations. In recent years, modeling has extended to dynamic networks [KO11], but some limitations remain. Consequently, one can take a step further and use temporal hypergraphs. Hypergraphs generalize graphs by extending edges to connect any number of vertices, allowing complex relationships to be described more accurately [SFE15] while reducing ambiguity and network inflation. Utilizing temporal hypergraph prediction models, however, introduces its own set of challenges.

First, as the model structure is more complex, it is relevant how the information is communicated to the analyst through visualization (cf. [HC14a]) and how domain knowledge feedback is incorporated. Static hypergraphs can be considered as standard sets, with different visualizations available [AMA⁺14]. Temporal hypergraphs, meanwhile, add a time-dependent evolution, making it harder to convey the relevant information meaningfully.

Secondly, many traditional graph-based concepts cannot directly be applied to hypergraphs. Hyperedges, as arbitrary sized sets of connected nodes, add another order of complexity. In previous works [AW18; ARW19b] it is shown how geometric deep learning can be applied to hypergraphs and showed how this method could be leveraged to predict behavioral patterns in social media hypergraph models.

Consequently, the incorporation of machine learning techniques into an interactive model to more accurately predict changes in the hypergraph due to changes in the data introduces new problems. While deep learning avoids assiduous manual feature engineering and algorithm design, it reduces explainability and accountability of the results. Domain experts usually have some domain-specific intuition—a mental model and structure—about inherent and implicit relations and groupings not available in the data, enabling them to judge the plausibility of hypotheses and to steer the exploration. Yet, they face difficulties articulating their domain knowledge through machine learning into the predictions and tracing its influence.

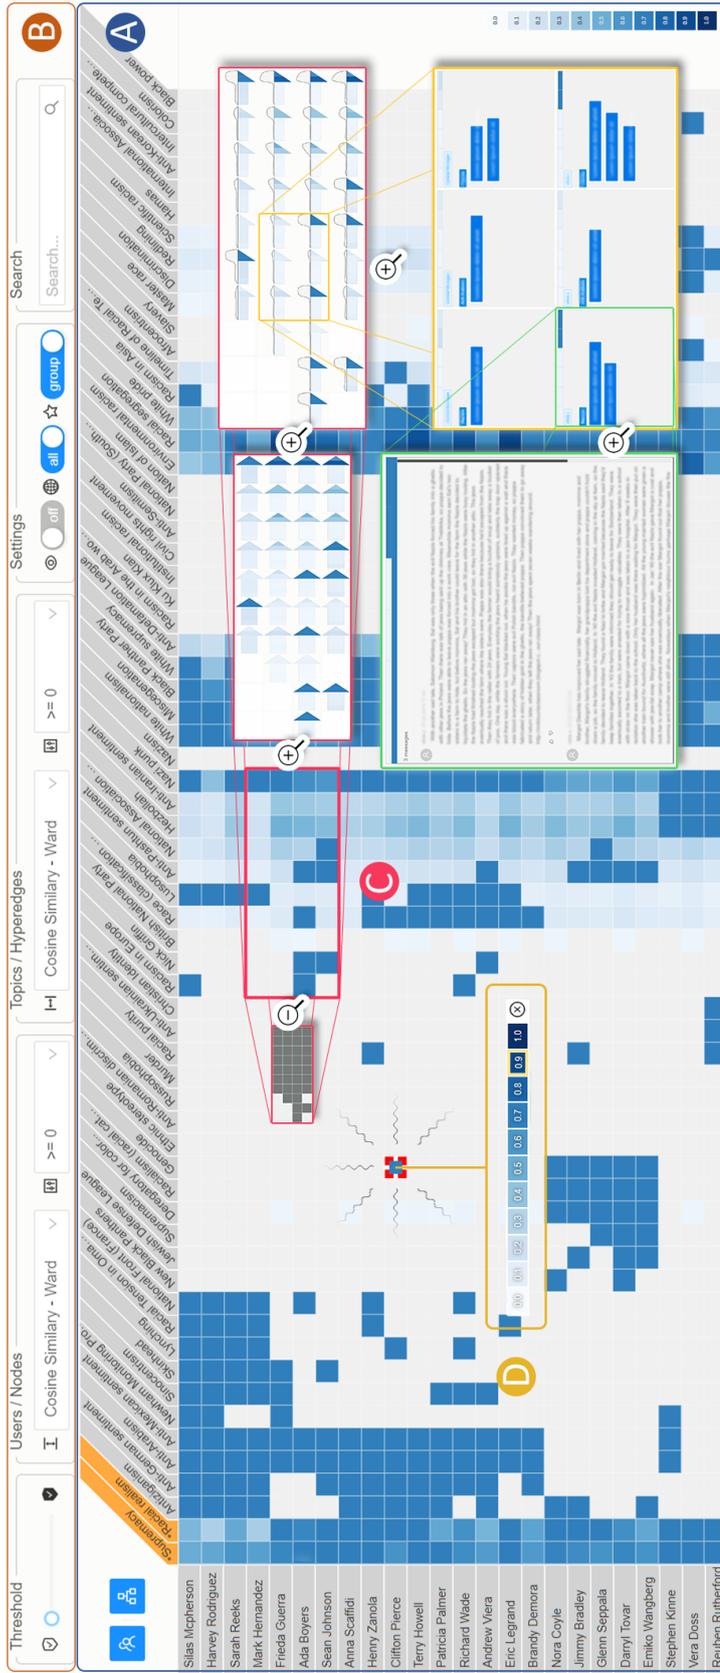


Figure 5.1: HYPER-MATRIX, a novel approach to explore and refine temporal hypergraph models using visual analytics. The interactive multi-level matrix-based visualization (A) enables the inspection of the model, together with the upper interface (B). The main area shows the second semantic zoom level applied to an obfuscated real-world dataset in criminal investigations, while the five insets (C) show the other drill-down levels for exploration. The technique allows to interactively (D) contribute domain knowledge, the resulting implications have ripple effects on the whole machine learning model, thereby refining it.

This holds especially for very complex models, like temporal hypergraphs. The knowledge formalization requires a very detailed a priori understanding of the problem by domain experts, which is not always available. For the same reason, it is challenging to capture the knowledge independently of the model without rapid, iterative feedback. Hence, the machine learning outcome often correlates strongly with the adequacy of the initial problem modeling and the quality of the training data, while domain expertise and domain knowledge are frequently not leveraged to their full potential.

To address these issues, we present HYPER-MATRIX, making the following contributions:

Contributions

- A novel, interactive **framework** for temporal hypergraph exploration through the use of semantic zooming relying on a multi-level matrix-based approach and various exploration concepts.
- The extension of a geometric machine learning architecture [AW18; ARW19b] with a **relevance feedback model**.
- A tight **coupling** between the visualization and the machine learning relevance feedback model for evaluation and seamless refinement, offering the integration of domain knowledge and making the corresponding model changes visually transparent.
- One **case study** describing an application of the technique to law enforcement.
- A **formative evaluation** with law enforcement experts using real-world communication data, demonstrating that our technique surpasses existing solutions and enables the effective and targeted analysis of large amounts of information.

Our approach bridges the gap between visual exploration and separate model training, allowing domain experts to enhance the machine learning predictions with implicit domain knowledge in the same step as evaluating and exploring the temporal hypergraph model predictions.

5.2 Related Work

This research is an entry into the interactive temporal hypergraph model exploration in the context of explainable support by machine learning. In the literature hypergraphs are studied from both a visualization as well as a machine learning perspective. In the following discussion, we adhere to the same distinction and relate

our work to the visualization of temporal hypergraphs as well as their application in machine learning.

5.2.1 Visualization of Hypergraphs

We first shortly discuss the situation for (static) hypergraphs as well as dynamic graphs, before looking at temporal hypergraphs. Hypergraphs can be considered as a **set of sets**. The survey on set visualizations by Alsallakh et al. [AMA⁺14] shows that several visualizations are applicable to hypergraphs. Hypergraphs are often drawn as regular graph networks or bipartite networks. When making their dimensionality explicit, they can be drawn as subsets—like Venn diagrams or radial sets—or in **node-link** form [VBW17], using colored hulls or other, specifically adapted approaches [MRS⁺13]. A third possibility is to use a **matrix-based** approach, which improves scalability [KLS07]. Subsets and node-link diagrams suffer from limited scalability, quickly leading to occlusion and clutter. Bound in the number of visual attributes they can employ, these techniques typically reach their constraints in the order of one or two dozens of hyperedges [VBP⁺19]. Further, they are difficult to extend with a temporal component, having already used up most visual attributes.

In comparison to set-based approaches, dynamic graphs change over time, leaving the choice [BBDW17] between employing animation or an additional timeline component. The former puts significant strain on the mental map when many connections change, while the latter is limited by the available screen space in the number of discrete timesteps it can show. The survey [BBDW17] also points out that node-link diagrams remain the most commonly used type of visualization. However, these approaches mostly lack the extendability to hypergraphs.

When studying temporal hypergraphs, the issues arising from the dimensionality and the temporal nature all build up. Indeed, there is almost no prior work on the visualization of temporal hypergraphs specifically. Two notable exceptions exist, which allow visualizing—but not modifying or refining—temporal hypergraphs: First, the recent works by Valdivia et al. [VBF17; VBP⁺18; VBP⁺19]. Their visualization approach is also shown later in Figure 5.6c as part of the case study. Second, the previous work by Streeb et al. [SAKW19] introduces an in-line visualization of the temporal evolution. Valdivia et al. begin to tackle the research gap by proposing PAOHvis, thereby claiming to provide the “first [...] highly readable representation of dynamic hypergraphs”. While this is a strong claim to make, the literature review showed a broad diversity between the approaches, but none—except the two mentioned above—is directly suitable for temporal hypergraph visualization, supporting this conclusion. Utilizing the previously discussed approaches as substitutes for a tailored visualization often does not adequately leverage the additional information available with temporal hypergraphs and does not address the tasks that come

with hypergraph topology and evolution. For those, we refer to Section 5.2.3. Shortcomings in existing approaches include, for example, Streeb et al. providing only the prediction abstraction level in their visual interface (cf. Level 3 in Section 5.4.1). Similarly, this is true for Valdivia, although they support coloring by a group. This can lead to information overload, as filtering using thresholds is the only way to reduce the information. In contrast, usage of semantic zoom enables an exploration of the complete hypergraph (cf. Section 5.4.1) without the need to preliminary apply filters while enabling tailored visualizations showing detailed information when focusing on different abstraction levels. Prominent examples of matrix-based visualizations are the Zoomable Adjacency Matrix Explorer [EDG⁺08] that enables users to zoom and pan with interactive performance from an overview to the most detailed views and the visual analysis system of Behrisch et al. [BDF⁺14]. It features a flexible semantic zoom to navigate through sets of matrices at different levels of detail. Further, both Streeb and Valdivia, only support sorting by weights and average (cf. size ordering in Section 5.4.2), compared to our default matrix-based sorting, improving cluster identification. Significantly, all existing approaches aim at analyzing a fixed hypergraph model. None focus on interactively working *with* the model and iteratively improving it (cf. Sections 5.3.2 and 5.4.3).

At last, while not strictly related to the research on temporal hypergraphs *per se*, we want to mention approaches that are, at least partly, similar to ours, and also conventional tools so far applied in practice. Here we concentrate on how hypergraph-like data is handled in the law enforcement field, relevant for the case study and the evaluation through domain experts (see also Sections 5.5 and 5.6). The visual analysis of communication data—but without any hypergraph visualization or a tunable model—is not novel and has been researched both from the analytical side [LZ15] as well as the visualization side [WLY⁺14]. Also, the idea of semantic zooming for matrix-like visualizations has been described previously [Ham03], however, in a different way and in the area of software management. Further, it was also described how an overlay magic lens [GSBO14] can be used instead of zooming, to keep the context and allow for faster search space traversal from locations far apart, which we partly employ for the partition hierarchy (Section 5.4.2). In practice, for the law enforcement field, we found that data which benefits from a hypergraph modeling, like communication patterns or process analysis, is prevalent, but not supported by any system. Gephi [BHJ09] is sometimes used, but analysts often prefer Pajek [BM98; BM02], as it supports larger networks. The most popular tool is IBM i2 Analyst’s Notebook’s [IBM20] graph component due to the prevalence and familiarity in this domain.

5.2.2 Machine Learning for Hypergraph Models

The concept of learning with hypergraphs for semi-supervised classification and clustering, with the aim of modeling high-order correlations was introduced by Zhou et al. [ZHS07]. The approach expands spectral clustering to hypergraphs by presenting a label propagation technique that reduces differences in labels among vertices that share a common hyperedge. Hwang et al. [HTKK08] further explored correlations among hyperedges, assuming that hyperedges with high correlation possess similar weights. Recent studies [BBL⁺17] have focused on the parametric learning of weights, employing the propagation of node features across hyperedges [FYZ⁺19; YNY⁺19].

Understanding communication patterns of users on social networking sites has created opportunities for richer studies of social interactions and better prediction of behavioral patterns. The prediction of links in hypergraphs has been a popular topic in social network analysis, in particular in the multimedia domain. This involves the prediction of metadata information, such as tags and groups, for social network entities like Flickr images [AW18], music recommendations for Last.fm users through leveraging network proximity information [BTC⁺10], and the prediction of higher-order links, such as tweets with specific hashtags, on Twitter [LXLS13]. Furthermore, hypergraph learning models have been employed to integrate complementary information from multiple modalities in multimodal data analysis effectively. Multi-hypergraph learning approaches have been suggested to handle incomplete multimodal data for disease diagnosis in neuroimaging [LGYS17], while frameworks have also been developed to learn a compact representation for each modality in a multimodal hypergraph using a tensor-based representation [ARW19a]. Although these studies have demonstrated the relevance of hypergraph-based learning in predicting implicit links within a social network, none of the existing approaches have established an interactive learning formulation that can incorporate user feedback as an external source of information to enhance the predictive capability of a model or even change the intrinsic properties, such as learnable parameters, of a model. Here, we extend previous work [ARW19b] on link prediction in communication networks capable of fine-tuning the trained model by incorporating external relevance feedbacks.

5.2.3 Tasks for Evaluation of Temporal Hypergraph Models

Tasks in temporal hypergraph analysis relate to dynamic networks and set comparisons. A task taxonomy of the former is provided in the survey by Beck et al. [BBDW17], and for the latter in the survey by Alsallakh et al. [AMA⁺16]. For temporal hypergraphs, in particular, the tasks sometimes substantially differ; for example, one being the analysis of changes of both connections and attributes over time.

The proposed technique does not directly fit with any existing task taxonomy, positioning itself between disciplines [AA06]. For a discussion on existing taxonomies and their applicability to temporal hypergraphs, we refer to the existing work by Valdivia et al. [VBP⁺19] and summarize only the main aspects here. Our technique supports not all traditional tasks in set analysis [AMA⁺16], and in dynamic network analysis [LPP⁺06; APS14; BPF14; KKC15], summarized in [BBDW17]. However, it provides support for several additional tasks relevant to our driving application. These include the clustering of related groups independently of their temporal connection, the inspection of shared attributes of connections, the following of temporal evolutions, while both retaining an overview and simultaneously being able to explore details. In short, the experts are interested in connectivity information involving both graph topology as well as attribute values, which can be separated between time ranges. One main requirement is the need to include external (domain) knowledge that is not directly available as raw data and includes conceptualized topics in line with their mental categorization. These tasks are not sufficiently described or supported by existing taxonomies, as they neglect the additional complexity incorporated by hypergraphs and the domain knowledge integration.

Given the sparse research in hypergraph visualization, it is unsurprising that there is no prior work on bridging both fields; this is the gap we aim to fill: offering a technique that addresses the shortcomings discussed above, enabling the exploration and refinement of hypergraph models using interactive visualization, closing the visual analytics loop.

5.3 Extension of Machine Learning to Hypergraphs

In the following two sections, we describe the overall workflow of our approach, shown in Figure 5.2. We begin with an exemplary description of one geometric deep learning model, adapted to a task relevant for our law enforcement domain experts: the temporal prediction and analysis of patterns in communication data. It acts as a blueprint for problem-specific temporal hypergraph models. In Section 5.4, we then discuss the interactive exploration using visual analytic principles.

5.3.1 Notation and Formulation of a Temporal Hypergraph

In set theory, an undirected hypergraph $H = (V, E)$ is defined as an ordered pair, where $V = \{v_1, \dots, v_n\}$ represents the n vertices (hypernodes) and subsets of these vertices $E = \{e_1, \dots, e_m\}$ constitute the m distinct hyperedges. H is represented by the incidence matrix $\mathbb{I} = |V| \times |E|$, with entries $i(v_j, e_k) = 1$ if $v_j \in e_k$ and 0 otherwise. We define the neighborhood of v_j as the set $\mathbb{N}(v_j)$ of nodes within the same hyperedge as v_j .

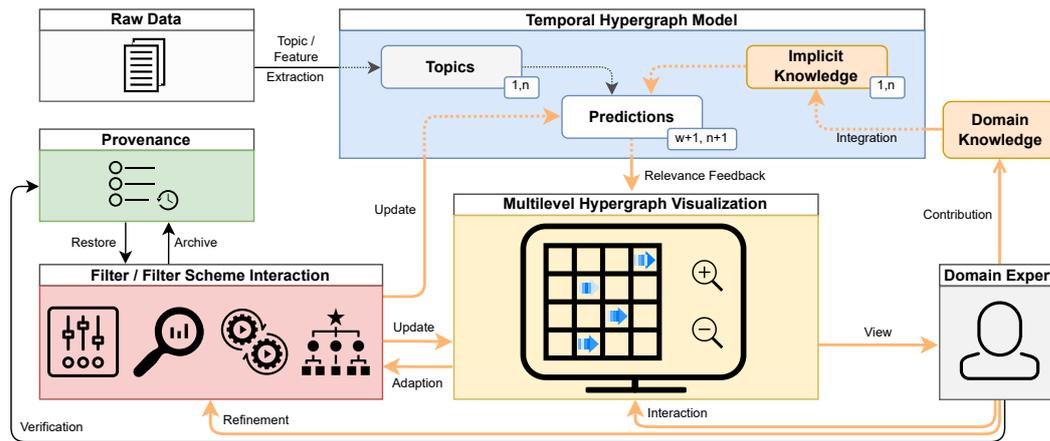


Figure 5.2: High-level workflow of our technique, showcasing the main components and the interaction flow for the exploration and refinement of temporal hypergraph models, adapted to use case A in Section 5.5. The workflow begins with **raw data** extraction and the generation of a **temporal hypergraph model**. The model state is visualized using a matrix-based **multilevel hypergraph visualization**, allowing for various exploration and **filter schemes**, including search and filters, a dynamically modifiable partition hierarchy, and matrix-reordering techniques. The **domain expert** can interact with the model by either refining the filter schemes or by contributing domain knowledge, which both update the model. The model feedback can then be explored and accepted, closing the **visual analytics loop**. The chronology of interactions and contributions are available for recovery or verification as a **provenance history**, facilitating accountability.

In adapting a generic temporal hypergraph model to our use case, we follow previous work [ARW19b], representing the relationship between internet forum users and their behavioral characteristics (both “explicit” and “implicit”). The available metadata (in particular forum category) forms the explicit characteristic of a user, while their topics of discussion outline the implicit communication characteristic. Thereby, we construct two separate hypergraphs depicting the connection of users with these explicit and implicit behavioral characteristics. To model the temporal component, let us define a temporal hypergraph by $H_{[t]}$, at a given time t , where each user is represented as a node, and each type of explicit/implicit characteristic is represented as a separate hyperedge. We denote the explicit and implicit hypergraphs, at any given time t , by $H_{[t]}^0$ and $H'_{[t]}$, respectively. Consequently, in $H'_{[t]}$, each topic is depicted as a separate hyperedge and users (nodes) who adhere to a common topic of interest are connected by it. Thus, forecasting the evolution of users’ topics of interest for time $t+1$ becomes equivalent to the task of finding new relations over the existing relations in hypergraph $H'_{[t]}$.

5.3.2 Relevance Feedback to the Deep Learning Model

Predicting links in temporal hypergraphs forms the basic idea for the model we use to forecast the interests of internet forum users. Performing link prediction on a hypergraph, denoted as $H_{[t]}$, with a fixed set of edges E targets to refine the set e_k . This link prediction task can be considered as a missing value imputation or a matrix completion task on \mathbb{I} , and can consequently be reformulated. In the following section, we extend previous work [AW18; ARW19b], thereby allowing to introduce feedback in HYPER-MATRIX. In the following we look at the mathematical formulation for training and then updating a geometric deep-learning model with user feedback.

Training Module Let $\mathbb{I}_{[t]}$ denote the incidence matrix of $H'_{[t]}$ at time t . This can be factorized as $\mathbb{I}_{[t]} = X_t Y_t^T$ with X_t and Y_t the row and column matrices, respectively. An hypergraph $H_{[t]}^0$ will serve as an auxiliary set of explicit information between users for predicting links in the implicit hypergraph $H'_{[t]}$. The Laplacian Δ_0 gives a measure for the relatedness between any two users [BTC+10]. The information in $H_{[t]}^0$ can subsequently be encoded by extracting its Laplacian. We can strongly enhance the user-topic link prediction outcomes by leveraging such a similarity measure, as it reduces noise and thus smooths the model output.

For model training, we employ a semi-supervised learning setup, hence the predictive loss is backpropagated by using a small set (around 5–8 %) of known links in $H'_{[t+1]}$. These known links create an upper bound for the number of timesteps the model can predict in $\hat{\mathbb{I}}_{[t+1]}$. Details can be found in [AW18; ARW19b]. Then, we take the incidence matrix $\mathbb{I}_{[t]}$ at time t and use the hypergraph link prediction model \mathbb{H}_{GDL} to learn the best parameter set $\Phi[t]$ for predicting the incidence matrix $\mathbb{I}_{[t+1]}$ at time $t+1$:

$$\hat{\mathbb{I}}_{[t+1], [t]} = \mathbb{H}_{GDL}(\mathbb{I}_{[t]}, 0) \quad (5.1)$$

Feedback Module We propose a novel interactive learning formulation to incorporate domain expert feedback into the underlying model in order to integrate their domain knowledge. This feedback is assumed to contain definitive implicit information about the topic of interest for certain users in the dataset. Rather than merely updating the information by changing the topic (hyperedge) of the respective users (nodes), the feedback should create a ripple effect on the overall connections in the hypergraph $H'_{[t]}$.

In more detail, if the feedback $f_{[t]}$ at time t involves the single user (u_j) denoted by node v_j in the hypergraph $H'_{[t]}$, then incorporating $f[t]$ will consist of a twofold operation: 1. *Update*: The topics for user u_j are refined, i.e., add/remove v_j to/from the respective hyperedges $E = \{e_1, \dots, e_m\}$ corresponding to $f_{[t]}$. 2. *Predict*: Adapt

topics for users in close communication with u_j based on their relatedness to u_j . This involves re-calculating the connection strength for vertices in $\mathbb{N}(v_j)$ with the hyperedges $E = \{e_1, \dots, e_m\}$. The first step is a straightforward updating of the matrix $\mathbb{I}_{[t+1]}$, which can be achieved by updating the values corresponding to nodes and edges suggested in the feedback $f_{[t]}$. The updated matrix $\mathbb{I}_{[t+1]} + f_{[t]}$ is then used as input to our link prediction model \mathbb{H}_{GDL} to consider the changes in the neighborhood connections. The feedback model does not follow an iterative process, but the learned parameters $\Phi_{[t]}$ are used as initialization for the model \mathbb{H}_{GDL} trained previously significantly faster after incorporating the feedback $f_{[t]}$ compared to starting anew. The representation of the feedback module in symbolic form can be formulated as:

$$\hat{\mathbb{I}}_{[t+1]} = \mathbb{H}_{\text{GDL}}(\mathbb{I}_{[t+1]} + f_{[t]}, 0, [t]) \quad (5.2)$$

5.4 Interactive Hypergraph Model Exploration

In this section, we focus on the visualization and interaction with the temporal hypergraph model, providing a tight coupling between the data manipulation and display (see Figure 5.2). We begin by describing how the model state can be depicted using a matrix-based visualization that provides drill-down capabilities across multiple levels via semantic zoom. Drill-down is thereby defined as the seamless zooming through the different levels during exploratory analysis, starting from a general overview to increasingly more focused and detailed information, as highlighted in Figure 5.3. To facilitate the interactive exploration, we present user-steering based on classical filters for standard search tasks, a dynamically modifiable partition hierarchy to include user-based structuring, and various matrix reordering techniques for the focused analysis of connections and groups. We then specify the interactions that allow domain knowledge to be incorporated into the machine learning model via relevance feedback and highlight how the updated predictions can be reflected in the existing visualization. This workflow facilitates the explainability of the underlying model, thus enabling the domain experts to provide more meaningful feedback. Finally, we describe how all interactions, domain knowledge input, and model output are stored in a provenance history, providing accountability and making the decision-making processes more transparent.

5.4.1 Model Visualization

As discussed above, the complexity of temporal hypergraphs makes them difficult to visualize. Hence, we propose a multi-level matrix-based approach, specifically tailored to the hyper-dimensionality as well as the temporal component. The visu-

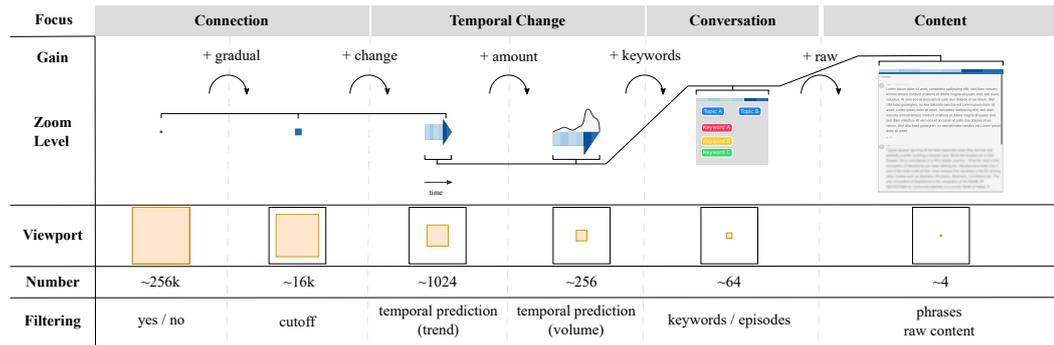


Figure 5.3: Semantic zoom levels and the different filtering levels as described in detail in Section 5.4.1. At each zoom step, the analyst gains another type of information about the model, filtering a different layer of complexity. As the focus becomes more detailed, the visualization takes up more space (zoom level and viewport as shown not to scale), while the number of visible entities decreases accordingly. The temporal predictions are shown in different forms throughout all levels (see fine grey line), with the detailed temporal evolution first shown in Level 3 and continuing down to Level 6.

alization (see Figure 5.1) consists of a menu bar on top, controlling the interaction concepts discussed later, and, for the main part, a matrix-like viewport, showing nodes as rows and hyperedges as columns, with corresponding row and column headers. This viewport provides freely pan-able and zoom-able drill-down capabilities across six levels of semantic zoom, shown in Figure 5.3, increasing or decreasing the information detail: from an overview of model predictions down to contents. For this purpose, we use three different level types: cells, arrows, and content boxes. Colored cell visualizations are used in Levels 1 and 2. An arrow-like representation reflecting a timeline is used in Levels 3 and 4. The base of the arrow represents the past, while the head reflects predictions. As the predictions become more uncertain with time, the arrowhead becomes smaller, reflecting the increased uncertainty and thus the decreased relevancy of the prediction. Levels 5 and 6 add text-based elements like keywords or raw content. Level 3 and beyond all contain the temporal aspect.

The visualization depends on the zoom state of the viewport. During drill-down, the focus shifts from a general structure overview over the temporal evolution to the raw content, providing the expert with more and more detailed information. Before we start with the description of this process, we define some necessary terms. As the feedback model outputs probabilities for the connections (see Section 5.3.2), gradual differences can be analyzed. When setting a minimum threshold for a connection to be meaningful, this allows for a *binary choice*. Showing a color encoding of the connection strength allows for a more expressive representation of the *gradual differences*. Setting a *cutoff* threshold can still be used to avoid cluttering with low-probability entries. The drill-down shifts the focus of the analysis. It starts

at the (binary) connectivity information, extends to gradual connection strength (Level 2), to the temporal change represented as an arrow (Level 3), to the temporal change encoded using position instead of only color (Level 4), then to information summarizing the underlying content for the predictions, in this case, keywords (Level 5), and, at last, to the raw data (Level 6). The design choice for an arrow glyph representation in Levels 3 and 4 is based on five reasons: (1) The principal idea of an arrow glyph was previously published [SAKW19] and found to be beneficial. Then, (2) given the target audience, a representation as an *arrow of time* is closely related to everyday experience. Further, (3) the separation into arrow base and head allows a clear distinction between past data and model predictions, which is very important for the target audience. The arrowhead also allows to visually reflect the decreasing prediction accuracy by becoming smaller. In terms of (4) visual advantages, an arrow provides a distinct shape, while, e.g., a cell is easily perceived to merge with neighboring cells, which is undesired. The choice also comes with disadvantages, introducing white space and can sometimes lead to distracting patterns. Finally, (5) a design study on combining timeline and graph visualization by Saraiya et al. [SLN05] shows that our approach—simultaneously overlaying the timeline—is best suited for detecting outliers. This is one of the main tasks for these levels, given the focus on change. The study also supports the design choice of showing only a single timestep in Levels 1 and 2, as the focus is on the topological structure. However, different visual representations like horizon graphs might be better suited when focusing on a continuous analysis. The seamless changes between levels speed up navigating through large models while eliminating demanding context switches. Moreover, at each step, the information becomes more complex, requiring more screen space to visualize. For a regular HD screen, we give rough guidance on the number of elements that can be usefully shown on-screen, amounting to around 256k grid cells of connectivity information and around four for the raw content.

5.4.2 Interactive Exploration and Drill-Down

To facilitate the interactive exploration, we contribute a user-steering based on classical concepts and filters for standard search tasks, a dynamically modifiable partition hierarchy to include user-controlled structuring and various matrix re-ordering techniques for the focused analysis of connections and groups. All these interactions concepts are reactive, and the visualization can smoothly and instantly update (< 100 ms), except for the domain knowledge integration in Section 5.4.3.

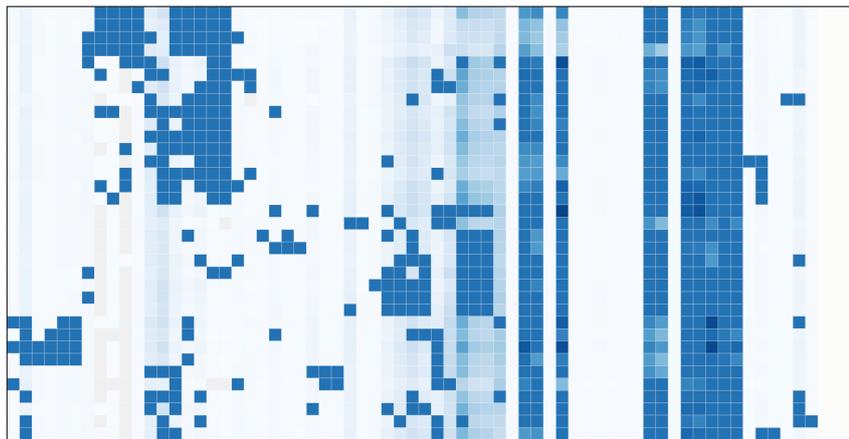
Interaction and Filter Concepts  Standard methods available in an interactive visualization are included, like (1) highlighting selected rows or columns, (2) highlighting hovered cells, (3) tooltip-based menus, (4) marking (i.e., starring) individual

entries to highlight them for tracking and follow-up, (5) adding textual notes, and (6) showing additional meta-information. Modal views allow to (7) control the partition hierarchy (see details below), while setting an (8) overall cutoff threshold allows controlling the confidence threshold of the underlying model. A (9) global search function provides the ability to search for node- and edge information as well as content and highlights the matching components. At last, the menu bar allows (10) controlling the matrix reordering (see detail below).

Dynamically Modifiable Partition Hierarchy  To allow domain experts to articulate their mental categorization to the model, the experts can create (nested) groups of different nodes or hyperedges, creating hierarchies. The nodes or hyperedges hereby relate to the leaves of the dendrogram. The groups can be expanded or contracted either directly from the node or hyperedge headers, visually indicated by color, or by editing them inside the partition hierarchy viewer in a modal overlay. The viewer shows a dendrogram-based representation with freely reorderable entries. Each branch of this dendrogram can be independently collapsed or expanded, i.e., the abstraction level is local to each branch and not globally set. For example, it is possible to collapse a large, uninteresting sub-branch, including the nested nodes it contains, while simultaneously having one branch fully expanded and another only up to the penultimate level. This is also independent of the overall visualization level, similar in concept to multiple fixed magic lenses, visually supporting different analysis paths. The hierarchy allows, for example, to group complementing entities together, to build meta-entities, and even hierarchies of entities.

Matrix Reordering and Sorting  To support the tasks relevant for our driving application (see Section 5.2.3), a matrix reordering is desirable such that related users and topics appear close to each other. Due to the independent and often conflicting interpretations of both axes and the sparseness of the underlying matrix, the direct application of standard 2D numeric sorting algorithms (e.g., Multi-scale-, Chen-, or Travelling salesman problem ordering) [BBH⁺16] often leads to unsatisfactory results, as they are mainly applicable to pairwise comparison matrices.

As part of the visualization, we offer three main different reordering strategies, as shown in Figure 5.4: (a) matrix-reordering (default), (b) sorted by size (connectivity), (c) first occurrence (original). The reordering is applied individually for each axis, as the requirement may differentiate between search tasks, not always favoring a block-like clustering. It also provides more flexibility for adopting other sorting methods in domain adaptations of our technique. The underlying sorting principles build upon a dendrogram-based serial matrix reordering discussed by Behrisch et al. [BBH⁺16]. It forms a multi-step process, combining the sorting of node and edge similarity vectors. Supported dendrogram methods are ward-, single-, average-,



(a) Multi-step hierarchy



(b) Size



(c) First occurrence

Figure 5.4: Comparison of different matrix reordering techniques to facilitate the detection of similar groups and connections. Compared to the unordered state and the slightly improved ordering by size, the adoption of a default multi-step, dendrogram-based reordering, modified and adapted from [BBH⁺16], enhances the clustering by similarity.

and complete linkage, combined with any pairwise distance function like Euclidean, cosine, or Jaccard. We refrain from discussing individual choices, which can vary strongly on domain adaption. For our case study, the Jaccard and cosine distance provide consistent results.

5.4.3 Visual Analytics for Model Updates

To increase the traceability of domain knowledge integration and explainability of the resulting model changes, we propose an interactive change feedback visualization, that seamlessly integrates with our visualization. The two-step process is shown in Figure 5.5. An expert can integrate domain knowledge by selecting a cell and setting a new connection strength (Figure 5.5a) and thereby complement missing or override model *input* data. This input is used to partly retrain the model and refine its predictions as described in Section 5.3.2, leading to a ripple effect. Thereby, the model has prediction authority, i.e., the user cannot manually fix the ultimate output to guarantee model authenticity. A spinner indicates the few seconds long operation. The resulting *changes* are displayed inside the same view (Figure 5.5b). A diverging color scale is used, showing changes instead of predictions. Through two visually distinct scales, it is immediately apparent if predictions or changes in the predictions are shown. The view integration allows for consistency, reducing the mental workload, and improving mental mapping.

Changes can be inspected on all levels of the visualization. The exploration is *not* restricted to just the current viewport, finding even weak connections. Change detection is facilitated, allowing rejection if deemed implausible or acceptance if convincing, enabling the followup of multiple analysis paths. By iteratively and interactively queering the model and see how it responds to domain knowledge integration, experts can discern better how connections and processes in the model are related, improving understanding and increasing explainability.

Experts in many applications are interested in their analytical progress and must reproducibly document the steps. We address this by a re-loadable provenance, storing the interaction sequence, domain knowledge input, model output, and fixed RNG seeds. This allows for inspection, verification, and traceability while providing accountability and making decision processes transparent. The provenance history allows undoing analysis steps, preventing dead-ends, revisiting and explaining past steps, but also bridging off to diverging analysis trails.

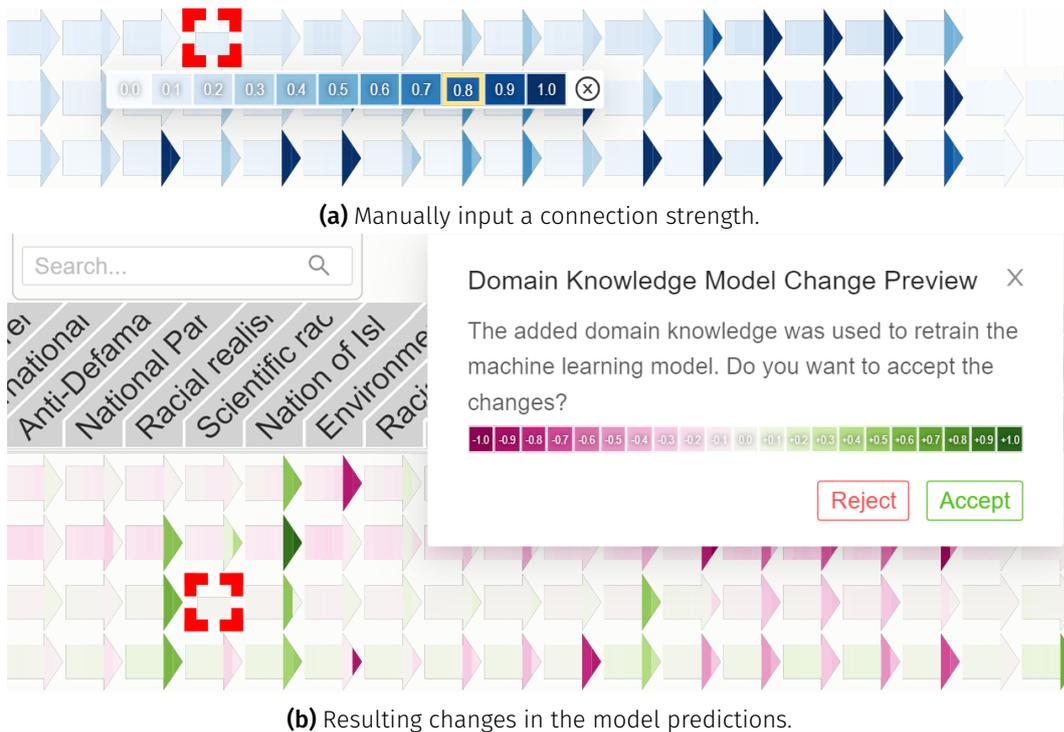


Figure 5.5: Resulting changes in the model prediction (ripple effect) from an **input** are visualized by a diverging color scale (from **negative** to **no** to **positive** change). They can be explored and **rejected** or **accepted**. This allows for model verification and multiple, different analysis paths.

5.5 Case Study: Forum Communication Data

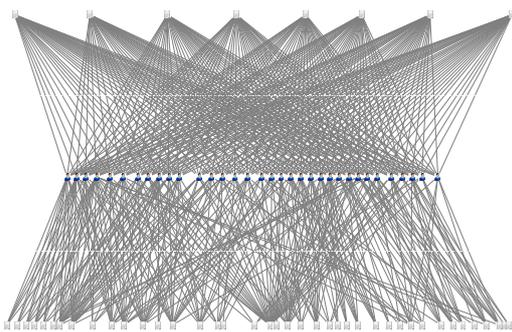
To demonstrate the visual exploration of temporal hypergraph models in HYPER-MATRIX, we conduct a case study, showing the applicability of our technique and improvements compared to existing approaches.

The communication data was collected from an internet forum well-known to law enforcement. It contains 335 188 text posts from 4904 users. We pre-processed the data using standard NLP methods to extract 158 topics, based on a domain-specific ontology. As described in Section 5.3, users are associated with nodes and topics become dynamic hyperedges. To allow for a reasonable side-by-side comparison with the existing approaches, shown in Figure 5.6, we had to restrict to a subset, consisting of 35 users, 65 topics, and six timesteps. This is around four times more than conventional approaches are designed for. We confirmed that our prototype works for significantly larger networks (cf. Section 5.7). Our prediction model is fed with four years (timesteps) of historical data and then predicts the evolution of the next two years as two timesteps. Almost any real-world data is noisy and may miss some relationships. Consequently, some of the conclusions drawn here may be

inaccurate. However, we focus on demonstrating the concepts and benefits of the visual analysis process HYPER-MATRIX provides.

The task we want to focus on in this case study is the **identification of related groups** and missing links, common in criminal investigations. To identify users discussing the same topics and topics discussed by the same group, the matrix reordering and connectivity information in Level 2 can be used to see structures, as shown in Figure 5.6b. Their spatial closeness acts as primary identification criterion, as similar row/column vectors are grouped closely. From this, their spatial closeness, describing the multi-step alignment, supports discovering related users or topics discussed simultaneously, but also latent connections. Distinct orderings can be applied separately to nodes and hyperedges, for example, to either favor overall similarity (cosine) or matching parts (Jaccard). For other requirements, it is also possible to include different metrics. To reduce noise and exclude weak connections, the top menu allows to set a threshold for the connection strength for historical and predicted data. A flag controls the ordering mode to either respect the filtered or the full dataset (including filtered elements). To further structure the view, the experts can manually click and select to group users and topics to reflect their mental categorization of users and topics. This allows to reflect domain-specific ontologies (e.g., similar concepts) or represent known formations of users.

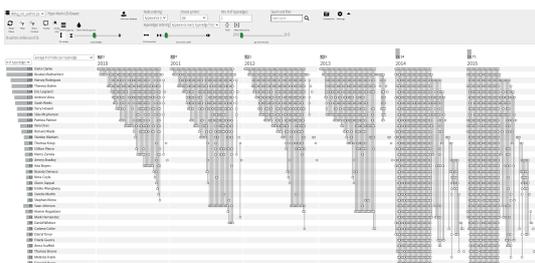
Zooming to the lower visualization levels shows the temporal development. Compared to existing approaches (see Figure 5.6c) our technique (Figure 5.6d) increases the scalability and comparability for dense temporal evolution. Compared to the industry standard Figure 5.6a, presenting the temporal evolution as a timeline-like arrow within each cell reduces comparison distances. Levels 5 and 6 allow an expert to understand the actual data on which a predicted connection is based: The main keywords of the relevant text fragments and, respectively, the actual raw text fragments (cf. Figure 5.3). This ability allows the expert to verify predictions and detect shortcomings as, for example, irony and coded synonyms are still difficult to be detected automatically. If the expert has identified shortcomings on any level, e.g., missing connections or wrong attribution of an ambiguous term, the technique allows for the inclusion of this additional domain knowledge. To externalize knowledge, the expert selects the corresponding connection and specifies the proposed strength on a scale between 0 and 1. This translates to definite knowledge about no and guaranteed connection, respectively. More nuanced values like .7 allow the experts to reflect their own uncertainty. This allows them to try out hunches while simultaneously preserving some model flexibility. For this reason, the change preview (cf. Figure 5.5b) is extremely relevant for the domain experts, as it allows them to see directly how their knowledge transforms the model prior to accepting the changes. They can explore the consequences by zooming and panning through all levels and correlate their findings with their intuition or other facts.



(a) IBM i2 Analyst’s Notebook. Automatically generated graph representation from the hypergraph model displaying the connections (labels removed) for the furthest predicted year using a modified bipartite representation. Data-wise, this can be compared to the connectivity information in our Levels 1 and 2. Clutter and occlusion prevent a meaningful global analysis, and while individual users and topics can be explored, this is slow, not without difficulty, and likely requires moving entities around to identify connections safely.



(b) Our technique at Level 2, showing the same predicted connectivity information as Analyst’s Notebook in Figure 5.6a. Clusters and related users/topics can be pinpointed more easily. The color scheme and filtering settings in the top menu bar also facilitate to identify the prediction strength, which can be estimated by using the overlaid legend in the bottom right corner. The blue buttons allow to access the partition hierarchy modifier to view a dendrogram view of the grouped entities.



(c) PAOHvis [VBP*18] shows the individual hyperedges, allowing to find connected users and topics. However, the hypergraph size is at the upper limit for a feasible visualization, already leading to some cluttering. Also, due to the temporal splitting, comparability between years is hindered for complex, non-sparse hyperedges compared to our technique, but better suited for comparing topics in the same year.



(d) Our technique at Level 3, showing the same temporal evolution information as PAOHvis in Figure 5.6c. The scalability is increased, showing no occlusion and the comparability of trends (important for the case study) is improved. This is due to retained cell ordering and short comparison distance. The downside is a reduced comparability between topics in the same year. The nature of the predictions is model-dependent.

Figure 5.6: Case study comparison of different approaches using the same internet forum hypergraph model dataset and exactly the same data view (connection strength > 0.1, min. 2 hyperedges). Compared are the state-of-the-art industry solution IBM i2 Analyst’s Notebook (Figure 5.6a), PAOHvis (Figure 5.6c) against our technique, showing the information at two different levels of abstraction (Figures 5.6b and 5.6d). Further, both external approaches only support a fixed network while our technique allows for an interactive refinement and domain knowledge integration.

If unsatisfied, they can go back. Otherwise, they can continue and repeat this visual analytics loop multiple times. This **rapid feedback** supports the expert in refining the model without being blind to the resulting consequences, but being able to control and explore the latest model state at all times. As the domain

experts focus is on exploratory analysis the iterative refinement supports finding connections and missing links faster. With domain knowledge that is difficult to be integrated a priori, step-by-step changes are more understandable.

5.6 Formative Evaluation

We performed formative evaluation sessions involving three domain experts (P1-P3). P1 is a criminal investigator working for a European law enforcement agency, having more than 30 years of experience, 20 years spent in digital and criminal investigations. His expertise includes communication and network analysis, familiarity with commercial systems like IBM i2 Analyst's Notebook [IBM20], the graph visualization tool Gephi [BHJ09], as well as the large network analyzer Pajek [BM98; BM02]. P2 works at the same agency in a different division, and has more than 20 years of experience in criminal investigations, specialized in group structure and content analysis. P3 is a senior project lead at a governmental research institute, studying analytical raw data analysis for more than ten years.

5.6.1 Study Procedure

The formative evaluation was conducted individually via remote screen sharing, taking about 60 minutes. For later review of these remote screen sharing sessions, they were recorded after receiving the formal consent of the experts. In the first 10 minutes a demo presented how to perform the visual analysis, explore and refine data and processes, and integrate domain knowledge in the search process and in the machine learning model. The next 30 minutes were spent between the experts using the system and providing feedback, as well as additional on-demand demonstrations. The tasks the experts performed include overview, the identification of the most promising leads, and the drill-down through the different zoom-levels down to the actual raw content, in this case, communication data. Further, we demonstrated and debated the different interaction techniques, like cutoff values and thresholds, matrix sorting and reordering strategies, and the dynamically modifiable partition hierarchy, as well as the machine learning feedback process.

In the last 20 minutes, the authors interviewed the experts asking 32 prepared questions. During each of the formative evaluation sessions, the experts engaged actively, trying out concepts, asking questions, commenting on the features, and pointing out issues. If an expert already partially gave comments during the 30 minutes session, they were offered to extend their answer. For example, when an aversion or surprising idea was mentioned, we additionally focused on these aspects. The interview was designed to elicit aspects of our technique that the

experts find relevant for their work or confusing or misinterpretable, as well as opinions on the individual approaches.

5.6.2 Findings and Lessons Learned

The **main observations** during the study are that our approach can effectively support most analytical requirements of the experts and that the experts favor both the **rapid exploration** of large datasets at different levels as well as the ability to integrate and contribute with their **domain knowledge**. This matches with their need to identify general trends in single combinations of users and topics and simultaneously identify co-occurrences. For this, the general prediction is more important than being able to identify differences between entities in the same year (cf. Figure 5.6). The underlying model we built upon [ARW19b] has proven to perform sufficiently well in this prediction task with an AUC (area under curve) of the ROC (receiver operating characteristic) of .88 and a recall value of .81. Excluded from the requirements are concepts outside the design scope, like purely mathematical capabilities as, for example, general centrality calculations, for which algorithms exist and could be included. In the following, we structure and summarize the main findings based on the expert's interactions and comments.

The domain experts agree that our approach of structuring information in **multiple levels** of details, using a **matrix-based approach**, is novel and therefore is not used in practice in their domain. For example, so far P1 has worked with either text-based or graph-based tools, and thinks our approach can “perfectly complement” existing workflows. The experts highlight the ability to effortlessly explore so much information (cf. P3), thereby “saving time” (P1), enabling a “quick analysis” (P3), while providing a “great overview ... with much details, ... but without overloading” (P2) the analyst, with an ease that is unexpected, given previous experience with this amount of data (cf. P2). We observed, that the experts often switch between the levels for targeting (upper levels) and then exploration and confirmation (lower levels). As P2 notes, this increases the size limit of the visually analyzable graph models, enhancing upon existing systems. “Together with the search capability” (P1), this allows for a very flexible workflow, enabling a good overview even for larger datasets.

The initial overview visualizations (Levels 1 and 2) are welcomed for providing a fast overview (cf. P1). The **color scheme** in Level 2 is regarded as comprehensible without explanation and aligning with expectations (cf. P1). It helps to provide guidance “where to start” (P1), and supports analysts in “planing their actions” (P3). To make the color scheme absolutely comparable, P3 requested the addition of a color legend. The **glyphs** are appreciated for providing details on the temporal distribution and future predictions (P2, P3). The glyph-based arrow representation in

Levels 3 and 4 is appreciated for providing details on the temporal distribution (cf. P2, P3) interesting to the experts, and, most importantly, “the future predictions” (P2) in context of the historical data. Depicting future predictions in the arrowhead and the past data in the shaft, and seeing both together was described as “helpful” (P3). The alignment by fixed timesteps, like years, is regarded as precise and practicable (cf. P1) by the experts. In comparison, the distribution as line chart in Level 4 received mixed responses, with P1 and P3 finding it beneficial for their understanding to get a better, absolute reading, while P2 feels “it does not add much”. The **keyword visualization** (Level 5) is regarded as fine for an abstract summary of the content visualization but could be extended (cf. P3). This layer, representing the “main connection” (P1) to the actual raw data, is important (cf. P1), and only shown when relevant in high zoom levels, “where the text content is relevant” (P1).

The ability to **search** through all underlying textual data and highlight matches in the views was received enthusiastically by all experts, as they can also transfer and fulfill some of their existing workflow, e.g., content- and text-based workflows, with our technique. It allows to explore global tendencies while enabling to query locally (cf. P2), not being distracted by other matches “not relevant at the moment” (P2).

While the visualization alone helps them already in some ways, providing them “with improved degree of detail ... unknown so far” (P1), all the experts also agree that the **interaction concepts** constitute an essential and relevant part of the approach, “helping them with strategical and operational decision” (P1). The **matrix reordering strategies** significantly improving the visual clarity of the overview, are regarded as “very interesting” (P2), and enable the experts to detect “groups” (P1) as well as connections easily, allowing them to “quickly identify hotspots” (P2), while putting less emphasis on weak connections. This is regarded as very supportive, being rarely supported in analysis systems (cf. P1), “saving costs and time” (P1). We observed that the experts use this as system guidance. The **partition hierarchy** is regarded by all experts as “essential” (P2), with P3 describing it as a “core functionality”. It allows grouping different model parts into physical concepts, applying structure comparable to existing mental models (cf. P2), improving the mental mapping. It “makes decision easier” (P3) and allows to “connect things” (P3).

The experts further describe that with existing tools, one major problem is that their **mental concepts** and models can “not [be integrated] enough” (P2) in the exploration, making it harder and less comprehensible. They notice that our approach supports them in three ways not present in existing tools: (1) the **interactive exploration** allowing to follow their instinct, (2) the modifiable **partition hierarchy** to express and capture their mental concepts, and, “most importantly” (P1), (3) the ability to integrate their **domain and external knowledge** directly in the model. While the experts wished that they could already “generate a report [...

and] export single entries” (P2) as commercial systems do, they note the enormous conceptual benefits of our technique. They regard them as “optimal” (P1), as there “are concepts and knowledge that cannot be modeled with machine learning [alone]” (P1) and are not “available” (P1) in the data. This knowledge then “cannot be integrated so far” (P1), is often documented in the head of the domain expert or “on a post-it note on the desk” (P1), leading to a high risk of the knowledge being “lost” (P1) or not leveraged. According to P1, the **knowledge integration** is performed iteratively during exploration, which we also observed as the experts adding knowledge intermittently, beginning with their main suspects and then expanding, adding knowledge when necessary either from post-its or when reading a name triggers a memory. The experts think that our **feedback loop** contributes to their analysis (cf. P1), replacing and “perfectly complementing” (P1) existing workflows. They regard the ability to *interactively* insert their knowledge as versatile. P1 noted that inserting all knowledge beforehand would be error-prone and “practically impossible” for larger datasets. To see “**validation** [possibilities] on changes” (P3) is especially important for vetting, and the change view is regarded as “very clear” (P1), allowing them a first glance, beneficial for prefiltering, steering and follow-up search guidance (cf. P1) to better divide their time for exploration. For improved usability P1 suggested to enable clicking to jump directly to the raw data in the change preview mode for validation. P1 regards the ability for a **global accept/reject** as sufficient for now, conceding that a partial accept could be explored in the future, although he does not see an immediate benefit. They state that the **0-1 scale** is “understandable and usable” (P1), but note that using the “5x5x5 system” (P1)—a commonly used police system based on letters A-E and 1-5 for source and intelligence evaluation [Nat10]—would be immediately understood and universally accepted in the target domain. The approach allows them to integrate their domain knowledge on multiple levels, together with the ability to perform a “quick analysis” (P3) of “large amounts of information” (P2) “in a targeted” (P1), non-overloading manner. From the observations of the experts, we derived a set of **tentative tasks**, relevant in law enforcement: (1) finding linked users/topics, (2) connecting users which share related topics to identify co-conspirators, (3) using classical text-based search in the raw data to identify users, (4) finding and judging an in/decrease of user activity for a topic, (5) finding a temporal co-occurrence between topics and users, (6) adding domain knowledge to a specific user and specific topic and judging the implications, (7) transfer raw data patterns and identify related users, and (8) confirming the model predictions by cross-validation plausibility with the raw data texts.

5.7 Discussion and Future Work

During the evaluation, we received multiple proposals on how our approach could be extended further, including by mathematical analysis methods and industry-grade interfaces. In the following, we discuss the limitations and broader applicability of our approach, also in the context of future work. For our prototype approach, we adapted the generic blueprint of a machine learning model to the case study. This use case has its own limitation, requiring structured data with time and author information, and dependent on advanced topic extraction models. We tested our prototype successfully with 1 000 users, 800 topics, and 15 timesteps on an HD screen, typically the upper size for large investigations. In terms of data type, the technique can cope both with sparse and none-sparse matrix structures. For the former, the matrix reordering allows to prioritize more relevant connections and order them further on the top left, reducing the required screen usage for the main parts. Of course, a *homogeneous* sparse matrix does not benefit from that. In this case, and for none-sparse matrices, the different zoom levels shift the size limitations. Nevertheless, they do not scale infinitely. Scrolling would be needed when scaling further, even for the overview level. According to domain expert P1, there the primary concern would be the number of users (y-axis), but using the partition hierarchy and matrix reordering could partially mitigate the issue. When increasing the number of time steps, the arrow becomes more detailed, shifting from blocks to a more continuous stream, becoming less distinguishable. For our use case, this fine-grained time is not primarily relevant because the experts aim at seeing who has recently been interested in a topic. However, it might become an issue when the task requires to extract detailed timestamps. Therefore one could use hovering, magnification on demand, or a more specialized visualization. Also, the visualization presented is better at analyzing trends and connectivity tasks on an overview level. Comparing the same time step in Levels 3 and beyond between two non-aligned nodes, however, becomes harder. For further work, we envision an adaptable overview layer showing a specific time point, allowing cross-cell comparability. When adapting to different use cases, some of the filtering methodology likely has to be changed. For example, when supporting biochemical process analysis, the raw attributes are not texts anymore, which (1) would need a different visualization for the content in the two lowest display level, but would also impact (2) the search functionality, which would need to be adapted to search and filter for biological and chemical properties instead of only text. The discussed visualization components serve only as examples for the visual analytics workflow presented. When adapting to a different field, there exist manifold possibilities for extensions, by integrating domain-specific visualization components. We provision this by a modular view architecture, supporting independent layer modules. Further

enhancements are multiple magic lenses to allow for simultaneous drill-down to different levels.

In the future, we envision improvements to the feedback system, for example, showing how domain knowledge propagates not only between two model states, i.e., before and after adding knowledge but also explaining the effects of previously introduced knowledge, for example, by interactively highlighting the individual influences on hover. This is supported by our architecture, but the computation time scales linearly with the number of domain knowledge inputs, which leads to computation times of several minutes and more, making it infeasible in an interactive environment for fast iterations. We hope to improve this by enhanced engineering, reducing the model setup and reloading times by advanced ways of updating the hypergraph model.

5.8 Conclusion

Many processes are difficult to describe using traditional graph-based concepts and benefit from more precise yet more complex modeling as temporal hypergraphs. We address this challenge by using a geometric deep learning approach and extend it to hypergraphs. However, such deep learning models typically do not incorporate domain knowledge, usually unavailable in the data. This is not least because domain experts struggle to articulate their knowledge without rapid, iterative feedback and intuitive representations matching their mental models, alternatively requiring a detailed a priori understanding of the problem. Hence, domain expertise is often not leveraged to its full potential.

We contribute a technique, named HYPER-MATRIX, to make temporal hypergraph model exploration more accessible for domain experts by enabling the integration of domain knowledge into the process and support their mental models through a multi-level matrix-based visualization architecture. The technique enables the interactive evaluation and seamless refinement of such models while providing a tight coupling and rapid, iterative feedback cycles to the underlying machine learning model. Model changes in response to the integration of domain knowledge are visualized transparently by a change preview, allowing experts to foster a more detailed understanding of how the underlying model works while externalizing their knowledge to teach the machine.

The approach allows to swiftly explore vast search spaces while maintaining focus and eliminating demanding context switches. Drill-down capabilities across multiple levels allow studying details and model contents on demand while retaining the overview. This architecture facilitates a focused analysis of relevant model aspects, allowing experts to detect patterns more rapidly and accurately. It is complemented by interactive filtering and search, various matrix reordering techniques, and a

dynamically modifiable partition hierarchy, allowing the integration of domain knowledge in the visualization layers.

We evaluate our approach in one case study and through formative evaluation with law enforcement experts using real-world communication data. The results show that our approach surpasses existing solutions in terms of scalability and applicability, enabling the incorporation of domain knowledge and allowing fast and targeted search-space traversal. While we focused on topic prediction for law enforcement as driving application, the interactions and concepts work with any temporal hypergraph, being model agnostic and applicable more generically to a wider variety of domains. With our technique, we pave the way for domain experts to a more interactive exploration and refinement of temporal hypergraph models, enabling them to use their knowledge not only for steering but also to articulate it into the machine learning model.

One cannot not communicate.

— Paul Watzlawick, *Philosopher*

6

Communication Pattern Interpretation through Conversational Dynamics

Large-scale interaction networks of human communication are often modeled as complex graph structures, obscuring temporal patterns within individual conversations. To facilitate the understanding of such conversational dynamics, episodes with low or high communication activity and breaks in communication need to be detected to interpret temporal interaction patterns. Traditional episode detection approaches strongly depend on the choice of parameters, such as window size or binning resolution. In this chapter, we present a novel technique for the identification of relevant episodes in bi-directional interaction sequences. We model communication as a continuous density function, allowing for a more robust segmentation into individual episodes and estimation of communication volume. Additionally, we define a tailored feature set to characterize conversational dynamics and enable a user-steered classification of communication behavior. We apply our technique to a real-world corpus of email data from a large European research institution. The results show that our technique allows users to effectively define, identify, and analyze relevant communication episodes.

Contents

6.1	Communication as Meta-Data Events	132
6.2	Related Work	133
6.3	Communication Behavior Modeling	134
6.4	Visual Analytics of Conversational Dynamics .	137
6.5	Conclusion	140

This chapter is based on the publication [SFS⁺19] and major parts of the following sections have appeared in:

- [SFS⁺19]: Daniel Seebacher, **Maximilian T. Fischer**, Rita Sevastjanova, Daniel A. Keim, and Mennatallah El-Assady. “Visual Analytics of Conversational Dynamics”. In: *EuroVis Workshop on Visual Analytics (EuroVA)*. ed. by Tatiana von Landesberger and Cagatay Turkey. EuroVA. Porto, Portugal: The Eurographics Association, 2019. ISBN: 978-3-03868-087-1. DOI: [10.2312/eurova.20191130](https://doi.org/10.2312/eurova.20191130).

For a statement of the scientific contributions, as well as the division of responsibilities and work in this publication, please refer to Chapters 1.2 (p. 22ff) and 1.3 (p. 26ff), respectively.

6.1 Communication as Meta-Data Events

With the digitization of society, especially in our daily communication, global information exchange has never been easier, resulting in mounting collections of communication data. The sheer amount, as well as the intertwined structures it is comprised of, pose challenging problems when trying to analyze communication dynamics. Questions such as—what are the patterns underlying the communication network or who are key players?

To address these questions, a variety of approaches were proposed, mainly, with a focus on social network analysis. Examples include the identification of key people in networks or the automatic detection of community structures [XSL11; XKS13; PBN17]. In the field of automatic text analysis, text content is examined more closely, for example using sentiment analysis [PL08], topic modeling [ESS⁺18], or lexical chaining [GRE15]. However, a problem that has not yet received enough attention is *how* people communicate with each other, i.e., a detailed exploration of the bi-directional interactions within a network. Such analysis allows to draw further conclusions about users’ behaviors and relations [EGA⁺16], thus allowing for more precise identification of roles in social networks.

In this chapter, we present a novel technique to support experts in their understanding of arbitrary, timestamped interactions, enabling a feature-driven investigation of relevant communication episodes. We use kernel density estimation to model the bi-directional communication events, based on their temporal distribution, as a continuous communication density function. In a second step, we present how to model features based on the communication density and other communication parameters which characterize the bi-directional communication

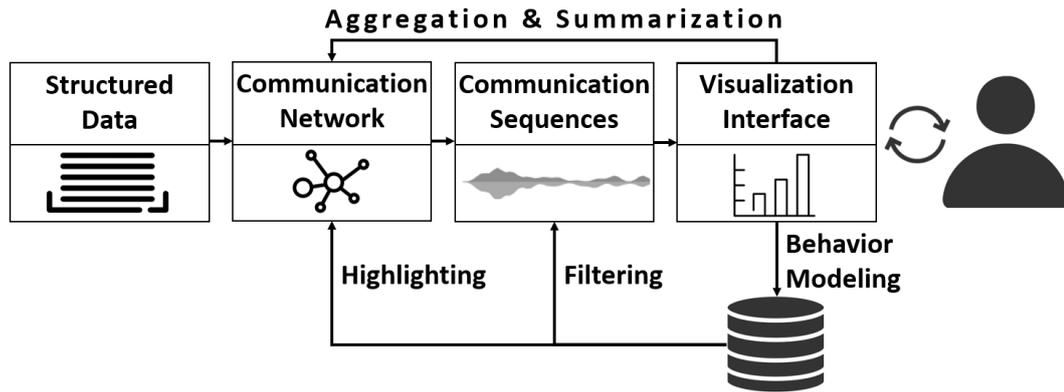


Figure 6.1: Visual analytics workflow of our approach for the user-driven investigation of conversational dynamics in large-scale communication networks.

behavior in individual episodes. By integrating this technique into a Visual Analytics workflow, as illustrated in Figure 6.1, we enable an investigation of communication episodes in large-scale communication networks that fulfills our identified tasks.

Overall, we make the following contributions in this chapter:

Contributions

- **(C1) A technique for modeling communication** based on the temporal distribution of communication events using kernel density estimation.
- **(C2) Density-based detection** of communication episodes in bi-directional communication sequences.
- **(C3)** Demonstration of how **features** can be defined and implemented to characterize the communication behavior in single communication episodes to allow for the visual analysis of those episodes.
- **(C4)** A **prototype** demonstrating the feasibility of this approach as a visual analytics approach for the investigation and analysis of conversational dynamics.

6.2 Related Work

Communication can be seen as social interactions involving numerous entities over time, which leads to large and complex networks. The task of analyzing such large networks is generally referred to as social network analysis, which is described in the standard literature [Sco17] and often focuses on using measures like centrality to analyze social ties and communication behavior [LZ15]. A general survey of visualization systems for networks is given by Shiravi et al. [SSG12]. Additionally, since such networks often contain the interactions of millions or billions of entities

over time, simplification is necessary, often using community detection algorithms such as SLPAw [XSL11] and CCME [PBN17]. An overview of other techniques is shown in the survey of Aggarwal and Wang [AW10].

Approaches that are related to our work and focus on analyzing relations and communications in graph networks include, for example, GestaltMatrix, a matrix-like representation [BN11]; TimeMatrix, which provides insight about the overall temporal evolution and the activity of nodes over time [YEL10]; *Timeline Edges*, which is an integrated approach and tries to leverage unused space in drawing zero-dimensional connectivity information as one-dimensional edges [Rei10]; *T-Cal*, a timeline-based approach that uses distortion to highlight areas with high communication volumes [FZC⁺18], or the methods proposed by Fu et al. recognizing communication patterns [FHN⁺07]. But all of these approaches have drawbacks regarding scaling, comparability, or information overload.

We also employ sequence analysis and, while the task itself is common, most approaches focus exclusively on statistical results or purely on visual comparison [MDM⁺15]. According to Zhao et al. [ZLD⁺15], only a few have investigated visualization approaches for comparing multiple event sequences. One idea that is proposed is CloudLines [KBK11]. Also, a metric has been presented for comparing temporal event sequences, but only for chains of sequences, instead of comparing sequences themselves [MDM⁺15].

6.3 Communication Behavior Modeling

For the analysis of the communication behavior, we concentrate primarily on the communications between an entity a and another entity b , for example, persons or communities. The communications between a and b can be considered as the multisets of the edges (a, b) and (b, a) in a communication graph. Different questions are of interest when analyzing the communication behavior between these two entities. For example, is the volume of communication high or low, is the communication discontinued, and is the communication one-sided (i.e., are there more communications from one entity to the other)? To answer such questions for a, b , we can compare the number of incoming messages from b with the number of outgoing messages from a , or vice versa. However, if we look at communications only as individual messages, it may be difficult to answer such questions. For example, for finding out if one entity is communicating more than another, we can compare the number of communications at a given time, but this is only possible if communications are compared for the same time ranges. If, for example, there is an hour difference between a communication from a to b and the response from b to a (which corresponds to normal response times for e-mails), this would only

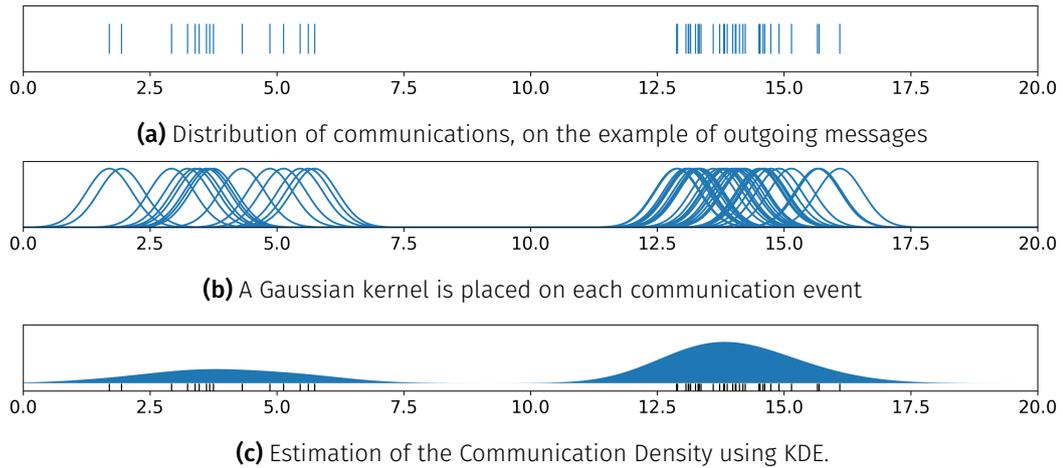


Figure 6.2: Individual communication events are represented as a communication density using kernel density estimation. The resulting continuous representation enables a robust detection of communication episodes, as well as, the derivation of features for a classification of such episodes.

be measured as a symmetric communication behavior if the communications were compared on the same time range.

In order to avoid these problems in the analysis of communication behavior, we do not model the communications as individual events, as shown in Figure 6.2a, but as a continuous communication density function, as shown in Figure 6.2c. This avoids the issues with binning or sliding window approaches as described above by using a smooth kernel. In turn, this prevents problems such as the failed comparison of communication behavior described above, since a communication no longer corresponds to a temporally atomic event, but can be measured with decreasing importance in the past and future and therefore no exact correspondence of the time units must exist anymore. In order to maintain this continuous communication density function, we use the well-known concept of KDE.

We replace every communication event between a and b by a Gaussian normal distribution:

$$G(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}, \quad (6.1)$$

as shown in Figure 6.2b, with μ being the center of the Gaussian kernel, i.e. the position of the communication event and σ the variance. We can then estimate the communication density \hat{f} for each time point x between the two entities, using the KDE:

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^n G\left(\frac{x-x_i}{h}\right), \quad (6.2)$$

with $h > 0$ as a smoothing parameter (bandwidth). The parameters μ , σ , and h can be adjusted as required to make this approach suitable for different domains and

tasks. The center μ is often set to zero (influence exactly around the event time), but could be used to encode a prior or subsequent response. The parameter σ describes the temporal influence an individual event has, where a very low value encodes a local event like the existence, whereas higher values could be used to encode more far-reaching concepts like a conversation about a specific topic, which continues for some time. The bandwidth parameter h describes how much individual responses like spikes should be retained, e.g. for occurrence of key words, or smoothed, e.g. for general tendencies. If we now consider the communications between two entities a and b , we can determine the communication density of the incoming messages $\hat{f}_{in}(a, b)$ (messages from b to a) and vice-versa the outgoing messages $\hat{f}_{out}(a, b)$.

By modeling communications as a continuous density function rather than as single atomic communication events, we can easily discover periods with a low or high communication density. For this, we can directly use the density functions f_{in} and f_{out} to judge whether one or both entities have made several communications in a given period of time. A further advantage of this approach is that it enables automatic detection of breaks in the communication (i.e., we can conversely identify individual communication episodes). For instance, very few people will continually send each other messages over long periods of time. Much more common is the pattern where one person sends a message that, in turn, leads to a discussion that ultimately ends after a few messages. We can determine these individual communication episodes by determining the periods s in which the communication density is greater than a threshold value.

Finally, to enable manual filtering of individual communication episodes as well as visual analysis, we demonstrate how a number of descriptive features for the analysis of communication episodes can be defined. With the help of additional variables such as the length L_{s_i} of one communication episode s_i and the density function for the incoming and the outgoing messages in this communication episode $\hat{f}_{out}^{s_i}$ and $\hat{f}_{in}^{s_i}$, we can then define features which are suitable for manual filtering and also enable a visual analysis of communication behavior of individual communication episodes. An example of such a feature would be synchronicity, i.e., if both entities are involved in a communication to the same extent at the same time. This would be illustrated by an equal communication density of incoming and outgoing messages in a communication episode. We can calculate this, for example, by determining the integral of the absolute difference between the two communication densities.

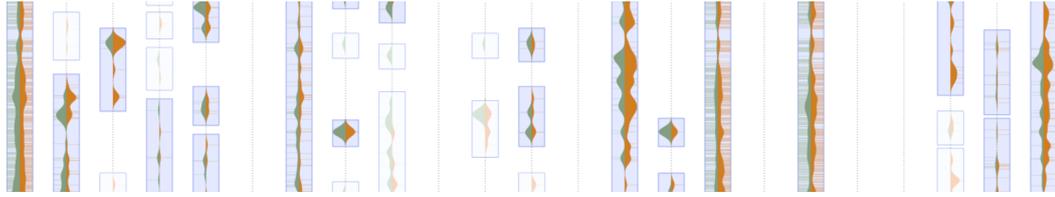


Figure 6.3: Communication Sequence visualization showing filtered communication **episodes** on a vertical timeline. **Incoming** and **outgoing** communication intensity is shown as a density distribution. In this case, episodes showing a strong *challenge-response* pattern are highlighted.

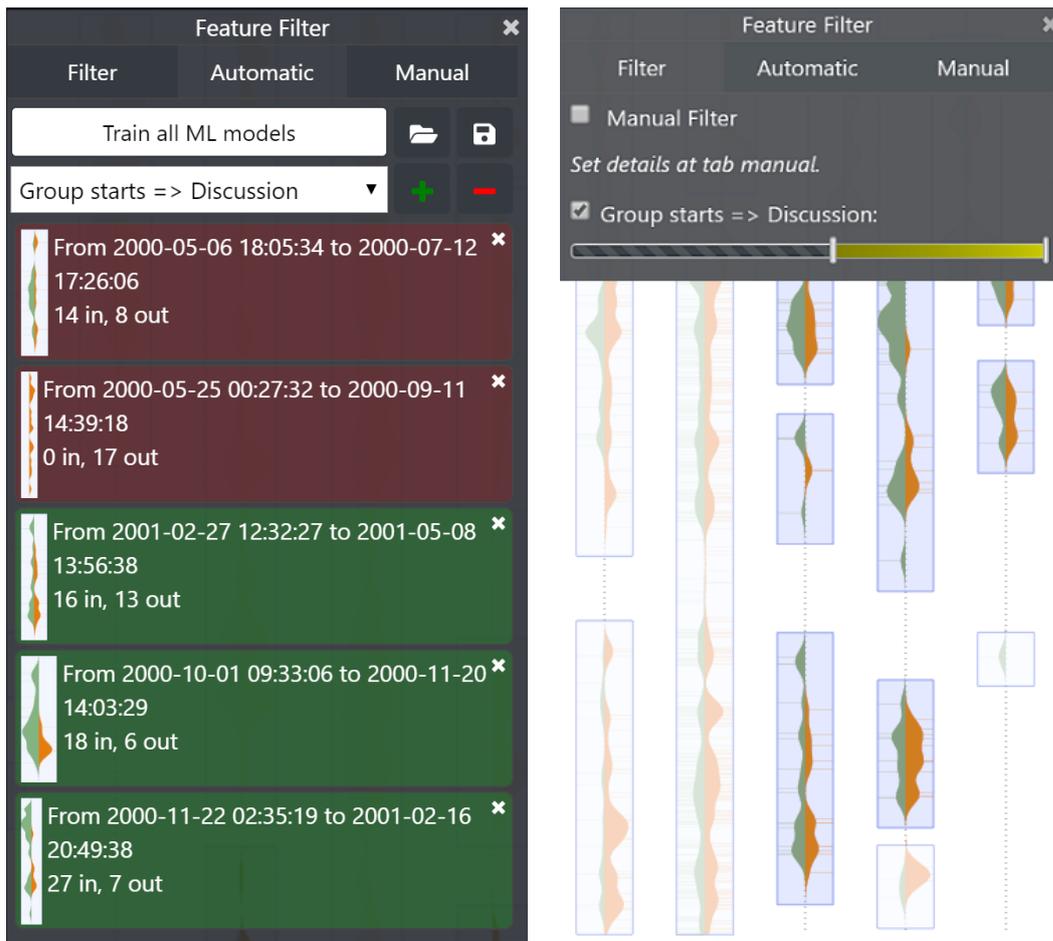
6.4 Visual Analytics of Conversational Dynamics

In the following section, we demonstrate how our technique, in combination with an experimental set of 14 descriptive features, facilitates visual analytics of conversational dynamics. To showcase a real-world dataset, we use email data from a large European research institution [PBL17]. The dataset is provided by the Stanford Network Analysis Project and contains the communication of 986 entities over a timespan of 803 days. In total there are 332,334 messages between 24,929 members of the institution.

Using communication **density**, we present a communication sequence visualization that enables identification of regions with low or high communication behavior. This communication sequence visualization also highlights the individual communication episodes. Finally, we introduce an interactive component that allows the user to manually filter the episodes as well as label existing episodes in order to perform a semi-automatic classification of the communication episodes into user-defined classes.

In order to look at the conversational dynamics in detail, we need to inspect the **temporal patterns** of incoming and outgoing messages more closely. To help with this, we have developed a visualization of the communication sequences between entities. To represent this conversational dynamic, we can use the communication density \hat{f}_i , defined above. We plot the density of incoming and outgoing communications \hat{f}_{in} and \hat{f}_{out} as area charts on different sides of a time axis, as shown in Figure 6.3.

For the visualization of the density of incoming and outgoing communications, we have selected the subdued colors lime-green and orange and optimized their contrast ratio. In addition, we can also use the communication densities to segment the communication into individual communication episodes by checking whether the density is above a certain threshold $\hat{f}_{in} + \hat{f}_{out} > \epsilon$. These individual communication episodes are highlighted to make them more distinct, for example with a light blue background. In order to visualize the conversational dynamics amongst multiple users, the individual communication sequences can be arranged side by



(a) Using **positive** and **negative** samples, a ML model is trained to identify episodes in which the selected groups start the conversation, leading to a discussion of both entities.

(b) Application of the trained model to the data. In this example only **relevant episodes** with high certainty are displayed, while irrelevant episodes are faded out.

Figure 6.4: By providing feedback for some data samples, users train ML models to identify relevant conversational dynamics in episodes.

side. In general, two arrangements are possible: (1) Vertical layout of the communication sequences, as shown in Figure 6.3, in order to leverage the width of the display to maximize the number of communication sequences shown. (2) Horizontal layout to leverage the width of the display to maximize the length of the shown communication sequences.

The concept of communication **episodes** also differs in their semantic relations, depending on the period under consideration. Communication encompassing several years has to be evaluated differently than one over several days. In the first example, messages may belong to the same episode, even though they might be several days apart. In the second example, however, this would be the entire monitoring period. It is therefore necessary to describe the high-level abstraction of communication differently, depending on the time range under consideration. These different concepts of episodes are supported in our interactive visualization

by semantic zooming. The available levels of granularity can be described by relative parameters, best adapted depending on the application domain and the specific analysis task, as described before.

To further enhance the comparability of the episodes, the concepts of timelines is extended; they can represent threads of time that do not need to be consecutive and can represent any number of time-ranges of an arbitrary length. Different pre-defined ranges like days, months, or, for instance, every Monday are available, while user-defined time periods are also configurable. If more than one linear timeline (the default) is selected, all timelines per group are juxtaposed. This makes it possible to compare the conversation dynamics at the same time in several years, which gives a better insight into recurring or changing communication dynamics. To provide further support, the whole view is interactive and each timeline is reorderable and realignable.

To allow for visual analytics of conversational dynamics, we need to be able to classify communication episodes into different classes. However, *a priori*, there is no predefined set of classes in which to classify the episodes. The desirable classes strongly depend on the domain and the analysis task under consideration. Therefore, we present a semi-interactive visual analytics approach where a user can define their own classes by example. A user can define a class and then provide some positive and negative examples as training data by clicking on relevant or irrelevant episodes. Classification is done using machine learning based on the defined features, which ideally show identifiable differences that reflect the user selection.

In our case, as shown in Figure 6.4, we use a Random Forest Classifier to make this binary match/no match classification with a confidence estimation since it can be trained with very few training samples. This trained classifier can be used to perform the binary classification for all other episodes, representing one model. It is possible to train several models and to combine them to allow for more advanced patterns. Theoretically, a completely manual approach can also work here, using rule-based classification. However, this becomes too tedious for more complex conversation classes and combinations of features and is therefore not practical. Using the semi-automatic approach, a user can define a class and train an appropriate classifier with only a few interactions. Since we use a Random Forest Classifier, we can model the uncertainty for the prediction of each episode. After a user has trained a classifier for a class, we can use this uncertainty measure to additionally filter the episodes. For example, the user can view relevant episodes for a class by choosing only those for which the classifier is very confident. In turn, this also means that we can inspect all episodes for which the classifier is very uncertain about the prediction. These borderline cases are the most promising

for re-labeling by the user in order to iteratively optimize the performance of the classifier.

Expert Feedback To evaluate the usefulness of our approach, we conducted an interview with one domain expert. For this interview, a different, proprietary communication dataset was used, whose characteristics are similar to the dataset presented here. The interview was designed as a combined system evaluation and feedback round. The following paragraph describes not only the key findings and comments by the experts, but also possible areas for improvement: The selection of non-consecutive, parallel timelines for comparability is regarded as useful, as well as the dynamic semantic zooming. Some fear was voiced that the default overflow of communication sequences to the right, to reduce the information density, might be misleading and lead to overlooked results. Therefore, it was recommended to compress the whole visualization on the screen initially—even when the density would be too high to be practical—and therefore require zooming all the time, but not leaving anything offscreen. The automatic detection of sequences with semantic zoom (levels of communication) in combination with filtering sequences and applying machine learning models to it is regarded as a very interesting, novel and realistic approach, which is useful to detect and replicate in other timelines or comparing between users. Both the manual filtering as well as the example-based machine learning are judged to be relevant, the former for first exploration and the later for comparison and detection. With these tools, the expert were able to semi-automatically find related patterns, which would be impractical manually.

In general, the expert interview showed the system works and that the approaches were received with interest and judged to be useful. According to the experts, the system offers many possibilities for different analysis tasks and is well suited for network exploration in the temporal analysis domain. Examples include the examination of bank transactions, phone records, or e-mails, where it proves very useful in specific situations, like finding relevant nodes. The main criticism voiced by the expert is the tendency for information overload when scaling the approach to show the conversational dynamics between numerous entities as they might occur in large communication networks, which might result in overlooked communication.

6.5 Conclusion

To demonstrate its feasibility, we applied our framework to parameters relating around communication density and response and have shown how we can visualize and analyze communication behavior with our modeling. This method, however, can be extended to encompass more complex domain-dependent concepts, for

instance, message content or sentiment. Apart from manually designed features, one can explore the emerging field of automated feature engineering as pioneered by Kanter and Veeramachaneki [KV15] and Katz et al. [KSS16]. Including own features enables a far more in-depth investigation of conversational dynamics. Nevertheless, the interview with the expert showed that our approach provides benefits when investigating conversational dynamics.

A challenging step for future work is to investigate how this approach can be used for the analysis of conversations of more than two parties, or how it can be integrated into a social network analysis workflow. A potential idea would be to use the communication episodes between entities, found with the help of our approach and classified as relevant by the user, for the weighting of the connection between the entities in a social network graph. Following our VA approach the user can also influence this weight by filtering non-relevant communication episodes. This weighting can then be used to steer community detection algorithms such as SLPaw or as an input for graph layout algorithms to visualize the social network structure. Thus, with the previously presented idea to include further domain-specific concepts, such as message content, community detection or layout algorithms could be further steered for answering questions such as whether discussions about relevant topics have taken place between users.

Part III | Holistic Approaches

Two souls, alas! reside within my breast, and each withdraws from, and repels, its brother.

— **Faust**, *Character in Faust* by J. W. von Goethe

7

CommAID: A Text-based Visual Analytics Technique

In the previous chapter, we studied the interpretation of patterns in communication primarily based on its meta-data. However, as we described at the beginning of this dissertation, communication can be more complex, consisting of both meta-information as well as content presented in a specific context. Currently, the automated analysis of communication data often focuses either on the network aspects via social network analysis or the content, utilizing methods from text mining. However, the first category of approaches does not leverage the rich content information, while the latter ignores the conversation environment and the temporal evolution, as evident in the meta-information. In contradiction to communication research, which stresses the importance of a holistic approach and which we discussed in [Chapter 2](#), both aspects are rarely applied simultaneously. Consequently, their combination has not yet received enough attention in automated analysis systems. In this chapter, we aim to address this challenge by discussing the difficulties and design decisions of such a path as well as contribute COMMaid, a blueprint for a holistic strategy for communication analysis. It features an integrated visual analytics design to analyze communication networks through dynamics modeling, semantic pattern retrieval, and a user-adaptable and problem-specific machine learning-based retrieval system. An interactive multi-level matrix-based visualization facilitates a focused analysis of both network and content using inline visuals

Contents

7.1	The Need for Holistic Communication Analysis	146
7.2	Related Work	148
7.3	Challenges and Design Decisions	151
7.4	Methodology: Modeling Communication . . .	152
7.5	System Design	154
7.6	Evaluation	162
7.7	Discussion and Future Work	166
7.8	Conclusion	167

supporting cross-checks and reducing context switches. We evaluate our approach in both a case study and through formative evaluation with eight law enforcement experts using a real-world communication corpus. Results show that our solution surpasses existing techniques in terms of integration level and applicability. With this contribution, we aim to pave the path for a more holistic approach to communication analysis.

This chapter is based on the publication [FSS⁺21] and major parts of the following sections have appeared in:

- [FSS⁺21]: **Maximilian T. Fischer**, Daniel Seebacher, Rita Sevastjanova, Daniel A. Keim, and Mennatallah El-Assady. “CommAID: Visual Analytics for Communication Analysis through Interactive Dynamics Modeling”. In: *Computer Graphics Forum* 40.3 (2021), pp. 25–36. ISSN: 01677055. DOI: [10.1111/cgf.14286](https://doi.org/10.1111/cgf.14286).

For a statement of the scientific contributions, as well as the division of responsibilities and work in this publication, please refer to Chapters 1.2 (p. 22ff) and 1.3 (p. 26ff), respectively.

7.1 The Need for Holistic Communication Analysis

The enormous growth in the use of electronic devices and systems in the past decades has led to an exponential increase in digital forms of communication. Simultaneously, the abundance of this digital communication [Sco09] and corresponding datasets has increased interest in how such communication can be analyzed in a wide variety of different domains, ranging from social sciences and digital humanities to engineering and business. For example, it has been studied how social and psychological features change with computer-mediated communication [GA08], how team performance can be assessed based on communication [FM08], how networks can be analyzed using text-mining [YP04], or how the evolution of dynamic communication networks can be visualized [Tri08]. This short list already shows a peculiar oddity when studying automated, digital communication analysis *systems*: most existing approaches focus on either the content of communication *or* on the network aspect—but not both. The first group of approaches usually leverages methods from natural language processing [MS99], while the latter uses techniques from the field of Social Network Analysis [Sco17].

Interestingly, this distinction between content and network is neither present in the seminal works on human communication research [Mor34; WBJ74], nor in modern works [Pea11] or current studies [FM08; Mes09], even if digital methods aid the manual analysis. Indeed, the analysis of network structure, communication patterns *as well as* content plays an integral part [SD60; WBJ74] of this research field. Abstractly, communication can be described as the exchange of meaning between entities, transporting information encoded into symbols [Pea11], reflecting the content's meaning and the network and metadata aspects through transport modalities. As such, analysis of the network/metadata or the content *alone* can sometimes provide a limited, incomplete, or even biased view on the communication, which is not ideal. Alternatively, employing independent approaches would introduce domain discontinues and complicates search tasks, placing an additional burden on the users.

The problem of how both the network and the content perspective can be combined has not yet received enough attention when considering communication analysis systems. This is especially relevant when such systems are used by non-communication experts, like in business intelligence applications or targeted criminal investigations, on which we based a case study (see Section 7.6.1). Typical tasks in these domains include searching for specific semantic content (e.g., negative product reviews, location names together with keywords), identifying groups (e.g., clusters or cliques), or a particular communication pattern (e.g., sequence, volume, timeline). These tasks can be addressed with separate solutions. But if, rather commonly, the search for specific semantic content needs to be restricted to a particular set of users (e.g., specific communication patterns, high centrality, or part of a specific clique), separate solutions struggle or even fail. Several other common tasks would benefit from such a combined search for cross matches and a fine-grained analysis of the communication network structure and context to detect such communication behavior and are therefore also not adequately supported so far. For example, detecting a terrorist attack on a soccer match by identifying the perpetrators increased chatter (hiding between the fans' increased chatter about the kick-off) about a parcel, which combines network, communication pattern, and content analysis.

In this work, we aim to address these shortcomings by discussing a possible technique as well as provide a framework for a holistic approach to interactive communication analysis. We do not aim to describe a turnkey system or replace existing solutions, but rather discuss the challenges and design decision in such a system, present an exemplary blueprint prototype on how such an integrated system could look like, and gather expert feedback on such a broader approach, to support further research and positively influence system development in this domain. For this, we extend upon two previous approaches presented previously, which we use

as building blocks. In [Chapter 6](#), we have used conversational dynamics to analyze communication patterns, covering the network analysis side. In [Chapter 5](#), we have described a technique for hypergraph analysis, combining machine learning and a multi-level matrix-based visual interface. From this, we borrow and adapt parts of the visual interface technique: The conversational dynamics aspects are integrated in this framework as one analysis model, while the multi-level matrix-based interface design is repurposed for the main interface.

In this work, we present COMMAID (Communication Analysis through Interactive Dynamics), making the following contributions:

Contributions

- A blueprint for a **novel, interactive framework** for a more holistic communication network analysis building upon individual models, providing a tight coupling between the network and the content analysis aspects.
- A description of **two extendable models** as example levels for network and content analysis, offering conversational dynamics [SFS⁺19] and semantic concept detection, also including several standard levels.
- A **discussion** on the challenges, design choices, and future work for a holistic communication approach.
- One **case study** and an **assessment** with eight law enforcement experts using real-world communication data describing an application in the law enforcement field.

Our approach bridges the gap between network and content analysis in automated communication analysis systems, supporting domain experts in exploring and analyzing arbitrary bi-directional communication. At the same time, we aim to pave the path for a more holistic approach to communication analysis.

7.2 Related Work

Social network analysis, which describes a collection of research methods for identifying structures in systems, is widely described in the standard literature [Sco17] and applied in many fields. For **communication analysis**, the early works were later extended and taxonomized by Bavelas [Bav50] and Leavitt [Lea51], before Roger [RK80] proposed to extend the field to communication networks. While detailed, domain specific content analysis, for example, in psychology [Ber52], was already known almost seven decades ago; only the recent advancement of computational capabilities allowed the focus to shift to a bulk analysis of communication data on a larger scale. Using methods from social network analysis, it became possible to

investigate **network aspects** like social ties and communication behavior [LZ15] by using centrality measures, detect communities [XKS13] and clusters [AW10] or model whole artificial networks [BMBL09]. However, using social network analysis on communication data primarily covers these network aspects. When focusing on **metadata and communication content**, a virtually unlimited amount of analysis methods can be applied. For example, metadata analysis [MHVB13] can be used to identify individuals, keyword-based searches [YP04] can filter for specific content, while methods from natural language processing [MS99] like sentiment analysis [GH11; PL08], topic modeling [RS10], or lexical chaining [GRE15] can support an advanced understanding of the meaning.

While this should give us ample scientific and technical methods at hand to analyze communication thoroughly, when we study automated, **digital human communication analysis systems**, we notice the peculiar oddity that most existing approaches focus either on the network or the content aspect, but not both, as we discussed above. A majority of the systems with communication analysis in mind follows the former approach. The de-facto standard in civil security and business intelligence applications are IBM's i2 Analyst's Notebook [IBM20] and Palantir's Intelligence Platform [Pal20], and, to a lesser degree, the large network analyzer Pajek [BM98], both **commercial solutions**. The open-source equivalent Gephi [BHJ09] is also used sometimes. While i2 Analyst's Notebook can be extended with content analysis capabilities, such a search is only offered as a separate interface. From a **visualization perspective**, all follow a node-link-diagram based approach. These suffer from inherent limitations like clutter or occlusion when the graph size becomes too large, and connections cannot be filtered enough for the search tasks, while techniques like edge bundling can only help so much. In fact, a study by Ghoniem et al. [GFC05] shows that matrix-based visualizations are better suited for large or dense networks and perform better from a scalability viewpoint. Various other methods are described in a general survey [SSG12] of visualization systems for large networks. For example, when considering communication networks as multivariate graphs, one could employ techniques like Multilevel Matrices [Ham03], LiveRAC [MMKN08], Hyper-Matrix [FAS⁺20] or Responsive Matrix Cells [HBS⁺20] for improved scalability and detail, often combined with matrix re-ordering techniques [BBH⁺16]. Alternatively, it is possible to leverage semantic or magic lenses [GSBO14] to highlight and enlarge relevant parts. When focusing on social network exploration, the survey by Riche et al. [HF10] focuses on specific extension for node-link and matrix-based approaches.

Looking at the **academic contributions**, we discover mostly alternative visualization and analysis methods. **Matrix-like** approaches are used, for example, by GestaltMatrix [BN11] to visually analyze asymmetric relations, or MatrixWave [ZLD⁺15] for comparing multiple event sequences. A notable set of approach that leverages

matrix designs were proposed by Nathalie Henry: MatrixExplorer [HF06] presents the idea of combining node-link and matrix approaches, which NodeTrix [HFM07] extends to address the occlusion problem for large node-link diagrams by switching to a matrix view locally. To address issues in path tracing in matrix views, they further present MatLink [HF07]. **Timeline-based** designs were proposed as part of Timeline Edges [Rei10] to efficiently use edge space, in T-Cal [FZC⁺18] to highlight areas with high communication volumes using distorted plot lines or as part of Cloud-Lines [KBK11] to display event episodes in multiple time-series. **Hybrid approaches** also exist, like Fu et al. [FHN⁺07] that propose to modify graph representations using multiple planes to recognizing communication patterns in e-mail networks. When considering the metadata and content analysis side, countless methods exist in various fields. However, many do not explicitly focus on communication analysis, and we will not discuss them here, although some can in principle be applied to a selection of the content-related tasks defined above (e.g., the interactive discourse analysis by Zhao et al. [ZCCB12]).

Leveraging analytical capabilities from both network and content information **simultaneously** has rarely been done. Interestingly, **commercial systems** seem to be ahead of their academic counterparts. Apart from Analyst's Notebook, which we discussed above, systems like Nuxi Discover and Nuxi Investigate [Nui20] for e-mail analysis and whole investigation frameworks like Palantir Gotham [Pal20] and more recently, DataWalk [Dat20] have become available. Some have received mixed responses by the public given their primary application in the intelligence and law enforcement community. As they are commercially developed, closed source solutions with few details about their detailed capabilities and internal workings, as well as their applications, they are often shrouded in secrecy (given their target domain). This proves problematic because it hinders oversight from an independent community like academia to track capabilities or point out issues like bias, which becomes increasingly relevant with the usage of machine learning techniques within these solutions. Looking at the **academic contributions**, we have three relevant approaches which combine network and content perspectives: TimeMatrix [YEL10] by Yi et al. combines meta-information and network structure. It uses a matrix-based visualization to analyze temporal social networks using TimeCells, showing a visual aggregate of a node's temporal information. For example, it can show edge count for a pair of nodes over a period of time. OpinionFlow [WLY⁺14] by Wu et al. combines content with network structure analysis to visually analyze opinion diffusion. They base their modeling on the network structure and sentiment-specific word embeddings, with the results shown using timeline-like visualizations. IEFAF [HDL⁺09] by Hadjidj et al. also combines content with network structure analysis. It uses a multiple-coordinated view system with a node-link diagram as the

primary visualization to support the forensic analysis of email, supporting various filter techniques, like metadata or keyword analysis and authorship attribution.

Given the little overlap between these solutions, their restricted applicability to communication in a generic case, and the growing support in commercial applications - in contrast to academic literature - reveals a missed opportunity. This is the **gap** we aim to fill: Provide a blueprint for a more holistic approach to communication analysis that supports network, metadata as well as content aspects simultaneously by the use of extendable plug-in models in a single interactive visualization system, enabling the effective exploration of communication for interrelated tasks.

7.3 Challenges and Design Decisions

Such an approach encounters several challenges. One obstacle comes in the form of the different requirements to internal data representation and the **analysis methods**, like graph-based approaches or content-based methods as well as the communication type involved, and how they can be combined in a single system while acting on the same data set. For this, we formalize communication modeling in Section 7.4 and describe the analysis as abstract operators working on a shared data space used internally.

The second challenge concerns the **visual representation and interaction** when combining these different methods in a single framework. The proposed system has to visually support different analysis modalities as part of a holistic framework through understandable and effective visualization methods, provide easy access to the visual results, and to allow useful interactions between them. The visual design choice also depends on the size and sparseness of the communication network. For example, a design based primarily on **node-link diagrams** like IEFAF might work well for very small networks, but requires larger ones to be sparse or decomposable into interrelated communities. However, choosing node-link diagrams makes it very hard to integrate additional information in an accessible way [Rei10], with coloring, overlays, and details on demand as options. An alternative [GFC05] is to use **matrix-based approaches** to support larger and denser networks, which also support in-cell content. We will follow this path for our approach. Compared to TimeMatrix, we extend the matrix-based approach considerably by using multiple views involving semantic zooming within the matrix visualization and thereby displaying specialized visualizations in-line and on-demand. Our approach scales well with the number of messages, which, as edges, are the primary source of clutter in a node-link-diagram. However, when the number of users exceeds several hundred [FAS⁺20], options like scrolling or magic lenses might be required. Our design focuses on a holistic approach, in contrast to existing approaches with limited exploration concepts or heavy task adaption. It offers flexibility when analysis tasks require combining

methods from different sub-fields of communication analysis. However, when analysis tasks are very specialized, for example exploring the network structure alone, a node-link approach might be more suitable. Further, a possibly viable addition to our design choice would be to use **coordinated views**, providing spatially separate visualizations that are logically linked. Such an approach could be explored further (see Section 7.7).

7.4 Methodology: Modeling Communication

In the following two sections, we describe the overall workflow, shown later in Figure 7.1. We begin by defining requirements for an abstract analysis level in Section 7.4.1 and define standard task levels to address common functionality. For the purpose of this work, a level can be thought of as a module which answers an individual analysis aspect. A level can have both a view (see Section 7.5.1) to show interactive visualizations and properties (see Section 7.5.2) to configure it. We give two exemplary descriptions of more complex, extendable levels as a blueprint for individual communication analysis. Firstly, a dynamics level in Section 7.4.3 to analyze network and metadata. Secondly, a level for semantic concepts described in Section 7.4.4. The architecture of the framework makes it easy to add own levels. While it is desirable for levels to cover distinct analysis aspects, they are not restricted from covering overlapping aspects. In Section 7.5, we then discuss the integration of the individual levels and the interactive exploration using visual analytic principles.

7.4.1 Abstract Level

A communication network can be described as a multidigraph $G := (V, M)$, with V a set of vertices representing the communication participants and M a multiset of ordered pairs of vertices representing a communication event. Additional metadata and content can be modeled by defining an information function $i : M \rightarrow D$, mapping a communication event to a data space D . Individual analysis levels can now be generically defined as operators that act on the vertex space V , edge space M , and the information function i . All this together forms the graph-like shared data space, which is used internally to store all information. Hashmaps and support index structures are used for efficient access. Each level can have none, one, or multiple in-line visualizations called views in the main interface (see Section 7.5.1). These visualizations can transfer domain- and task-specific information relevant to a domain expert. Further, each level has its parameters and filters for control. As individual, separate levels itself would not provide many benefits, the key idea is to complement each other on the system scale. Their flexible and simultaneous

application in a single approach provides support for cross-matches, as level-specific filtering adds together to form a global filter. Additionally, all levels can output a feature vector that is fed to a machine-learning-based retrieval system, described in Section 7.5.3, to enable intelligent user steering. The system can be customized and extended to more specialized tasks by adding additional levels to cover specific needs.

7.4.2 Standard Task Levels

Analysis usually requires a set of standard operators for filtering and selection, so we provide a set of standard task levels. For example, to support simple tasks like restricting the time ranges, one can define an operator on M which filters edges based on the timestamp information in the data space D (a **timefilter level**). Other examples are to filter participants in V through properties in M and D (a user **selection level**), or keyword-based search by restricting based on content information in D (a **keyword search level**). As these levels act primarily as filters, a corresponding view (see Section 7.5.1) might not be required. To provide basic visual analysis, one can define an operator on V and M (a **volume level**) which tracks the amount of communication between users, or an operator on V , M , and D , that track the temporal evolution of such communication (a **distribution level**). For both, we provide corresponding views in the main interface (see Section 7.5.1). For the remainder of this work, we will primarily focus on the more complex levels in the following two sections as they allow us to define task-specific analyses and only consider these standard task levels when necessary.

7.4.3 Dynamics Level

Different questions are of interest when analyzing the communication behavior between entities: For example, how does the volume of communication develop? Is communication discontinued? Is it one-sided, or are there specific patterns? However, if we look at communications only as individual messages, it may be difficult to answer such questions. To analyze the dynamics of communication events more thoroughly, we follow our previous work [SFS⁺19], and define a set of features which operate on the edge space M and the information function i . There, we model communication not only as individual events but as a flow, which can be described using distributions and a continuous density function. This view enables us to easily model influence effects like the response time (both prior and delay), the width of the temporal influence, or control between spikes and general tendencies simply by adapting parameters like μ , σ , or h , respectively. For details, see the original publication. It further allows us to detect breaks in communication and thereby identify *communication episodes*. The choice for the granularity of

the episodes is made globally and dataset-dependent. Together, this facilitates the structural communication analysis and helps to address the questions posed above.

In our work, we leverage these previous ideas and adapt them to work on top of the abstract level formalization defined above: We take the properties and visualization ideas from the original work and transform as well adapt them to work with this approach by adding level properties and an in-line visualization (a view), which we further describe in Section 7.5.1.

7.4.4 Thematic Level

In general, the inclusion of thematic concepts allows a user to refine their search task in a more powerful way than keyword lists and comes more naturally to analysts, who often think in concepts. Regarding modeling, a thematic level operates on the edge space M , depending on the content information in the data space D . A standard method to extract concepts from text data is named entity recognition. There, it is possible to either use pre-trained models or adapt them with domain-specific or task-specific concepts. However, a simple search using these semantic concepts might not be flexible enough to allow for more complex search tasks like "retrieve communication talking about a person in connection with a location". Therefore, we propose an interactive visual query language that allows for a flexible combination of semantic concepts to fulfill such search tasks. This query language allows creating multiple semantic queries based on spatial co-occurrence of semantic concepts. These queries can then be combined using Boolean algebra to build more complex filters. For example, the above search could be restricted further by additionally requiring an organization to be mentioned, which is combined with the first query. As this level acts more like a filter on the data, it is an example of a view-less level, not having a separate in-line visualization in Section 7.5.1.

7.5 System Design

This section focuses on the visualization and interaction concepts to integrate multiple levels in a single framework while providing a tight coupling between the network and content analysis aspects. The proposed workflow for this architecture is described in Figure 7.1.

We begin by describing how the overall network can be visualized using a matrix-based visualization that provides multiple levels as views, representing the individual analysis levels' results. Table 7.1 shows the interplay between Levels, Views, and their properties.

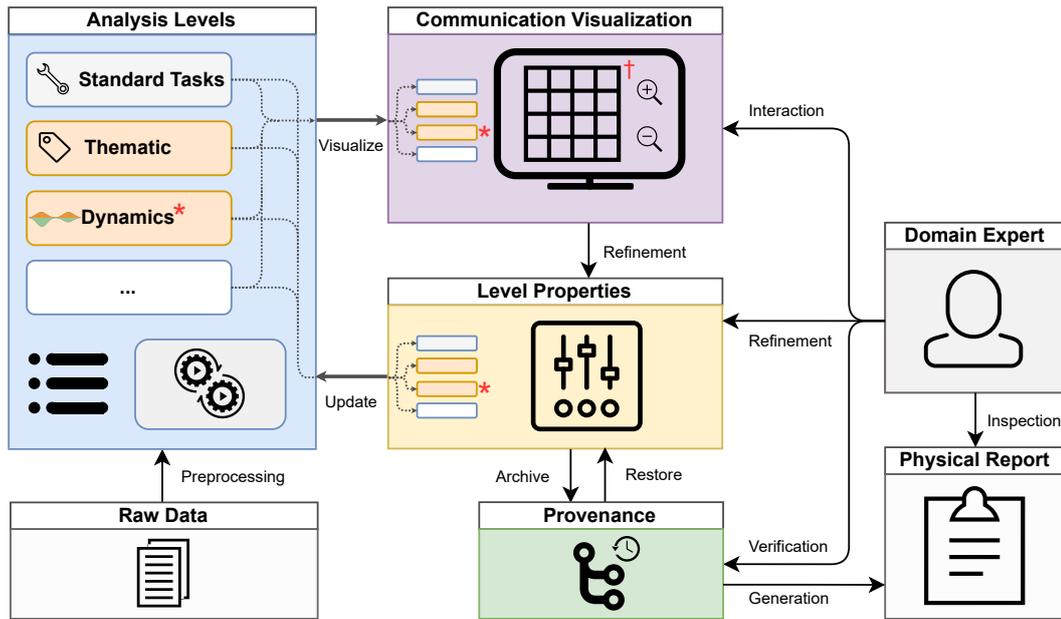


Figure 7.1: High-level workflow of our system, highlighting the main components and the interaction flow for the communication network analysis. The workflow begins with **raw data** extraction and the generation of the individual **level**. A **multi-level matrix visualization** containing other in-line visualizations presents the current model state to the user and allows for different interaction and exploration schemes. The **domain expert** can then explore and refine the levels through their **properties**, leveraging an internal **relevance feedback system**, updating the overall model state, and adapting the selection. The history of refinement interactions is archived to provide **provenance information**, which can be exported as part of a physical **report** for inspection, traceability, and explainability. Components indicated with a * [SFS⁺19] and † [FAS⁺20] are based on and extended from previous work, presented in Chapter 6 and Chapter 5.

Conceptually, the information becomes more nuanced during level drill-down, going from overview to specific analysis to content, while each level addresses a specific question related to the level-modeling in Section 7.4. To facilitate the interactive exploration, the levels can be controlled via a property pane. The levels can then act as filter methods, enabled through standard operators for standard task levels, steering options for conversational dynamics, and a visual query interface for thematic searches. We specify how these individual levels can act together, which in-line visualizations they provide to support the exploratory analysis, and how user feedback for the adaptable retrieval system can support the search. This methodology helps domain experts to gain a better understanding of the communication data by providing rapid-feedback through interactive filtering, covering different analysis levels simultaneously. Finally, we describe how all steps are recorded in a provenance history graph, making the decision-making processes traceable.

Level	View	Properties
Volume	X	-
Distribution	X	-
Timefilter 	-	X
User Selection 	-	X
Keyword Search 	-	X
Thematic 	-	X
Dynamics* 	X	X

Table 7.1: The available analysis levels in our system. Among the standard tasks levels, the two examples of more complex, custom level implementations are highlighted in bold. The component indicated with * [SFS⁺19] are described in Chapter 6.

7.5.1 CommAID Interface Design

For the visualization of the communication networks, we adapt the multi-level matrix technique  from our previous work [FAS⁺20]. However, we change the meaning of grid layout as well as the levels and the cell information: (1) Instead of representing nodes (rows) and hyperedges (column), they, respectively, become senders and receiver. (2) Instead of displaying increasingly detailed cell information, the in-line visualization represents the results of the individual levels discussed in Section 7.4 as independent views, which are different, but not necessarily more detailed. The overall interface of our approach is shown in Figure 7.2.

The novelty here is how existing visualization are combined, adapted, and integrated into an holistic framework and how the interaction with it is designed. Apart from the interactive matrix-based visualization **A**, the linked level property pane **B** allows to restricting the search space, using standard task filters, dynamics settings, and a thematic concept builder **C**. In this prototype, three different views are provided: *Volume* **D**, *Distribution* **E**, and *Dynamics* **F**. They are shown as part of Figure 7.2 and in more detail for three generic cells each in Figure 7.3. A provenance history graph **G**, discussed later in Section 7.5.4, allows to keep track of the analysis steps and results. In the following, we explain the design rationale of the views. It is important to note that through semantic zoom, the order in which the views are shown is fixed.

The basic principle of semantic zoom is that each cell of the matrix visualization  serves as a canvas for a different type of analysis result of the communication between two entities in the network. However, rendering detailed visualizations there makes sense if a cell has a specific minimum window size. Otherwise, even basic visualizations can be impossible to read. Guidelines [FIBK17] have been developed for the required size to retain readability. Along those lines, and with the type of views in mind, we have chosen a view switch with every doubling of cell size. When using a different type of view, the transition criteria might have to be adapted.

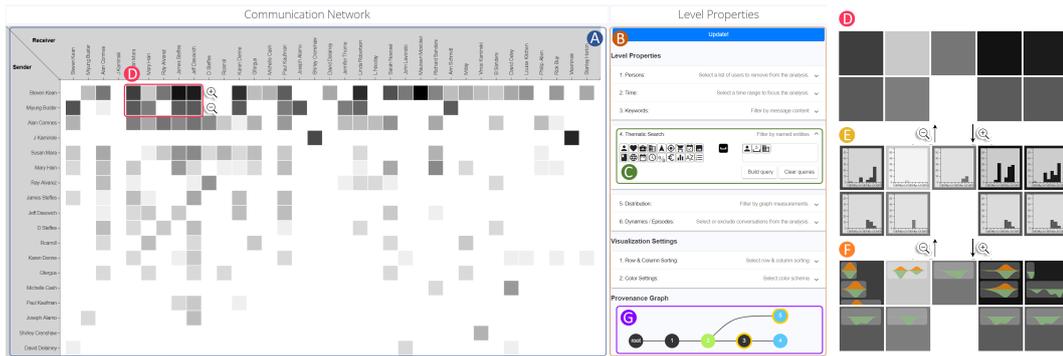
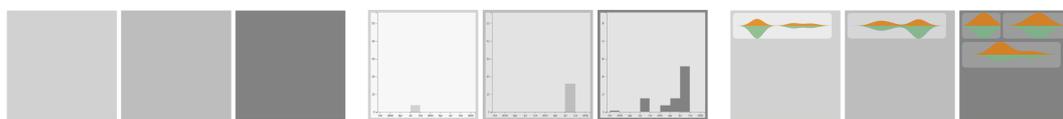


Figure 7.2: COMMAID, an integrated visual analytics technique to analyze communication networks through dynamics modeling, semantic pattern retrieval, and an interactive multi-level matrix-based visualization **A**. This visualization enables the inspection of individual communication at different analysis levels through semantic zooming, while the linked level property pane **B** allows to restricting the search space, using standard task filters, dynamics settings, and a thematic concept builder **C**. The in-line views show the visualizations provided by the individual analysis levels (different zoom steps) presenting volume information **D**, statistical distribution information **E** as well as communication episodes using conversational dynamics **F**. A provenance history graph **G** allows to keep track of the analysis steps and results. The technique allows to interactively explore communication activity from a metadata (connectivity, closeness, time, statistics) as well as a content level (keywords, thematic concepts) simultaneously, reducing discontinuities.

For example, either by using a different scaling factor or keeping the cell size for some view transitions and just switching the view.

The **Volume View** (belonging to the volume level) displays the number of communications between two entities, where the amount of communication is visually mapped to the cell's color. Different color scales can be used depending on the task requirements. Figure 7.3a shows a sequential, single-hue gray color scale, where white indicates that no communications are taking place, and black represents the maximum number of communications between two entities in the network. Color schemes are replaceable, for example, for users with visual impairments or by using diverging color schemes to indicate deviations from the average.



(a) Volume View: The amount of communication is visually mapped to the color of the cell. **(b) Distribution View:** Provides an overview of the temporal distribution using barcharts. **(c) Dynamics View:** Visualizes the communication episodes between two entities.

Figure 7.3: Overview of the three views that are provided in COMMAID. We distinguish between the *volume*, *distribution*, and *dynamics* view. The first two in-line visualizations come from standard task levels (see Section 74.2), the latter from the dynamics level (Section 74.3).

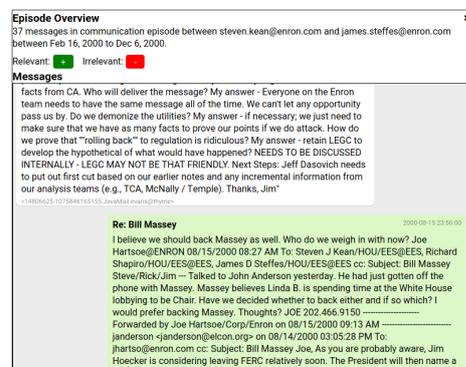
The **Distribution View** (belonging to the distribution level) is used to provide an overview of the temporal distribution of communications. Similar to TimeMatrix, we use a bar chart, but add a background color (matching the Volume View) to retain the overall amount information. Thus, in addition to the temporal distribution of communications, global information can also be visualized.

While these two represent views for the standard task levels, providing views for the custom models is especially interesting. Here we offer the **Dynamics View**, visualizing results from the dynamics level . There, we represent the communication episodes between two entities in the network. Depending on the tasks, the episodes can be shown chronologically or customly sorted.

All three views have in common that they offer additional details-on-demand. A click on a cell opens a zoom-level-dependent tooltip (see Figure 7.4a), which provides information about the time distribution, named entities used, or raw data. A click on an episode also opens a tooltip (see Figure 7.4b), which visualizes the discussion content between two entities using a chat-style metaphor. In both details-on-demand visualizations, the user can directly perform a refinement step, e.g., by excluding entities from the search or evaluating communication episodes for relevance.



(a) Details-on-demand populated by the distribution and thematic level presenting the time distribution, thematic named entities used, and, additionally, the raw data.



(b) Details-on-demand for a communication episode visualizes the discussion in a chat-style metaphor. The communication can be ranked for the classifier discussed in Section 7.5.3.

Figure 7.4: Overview of the details-on-demand offered by different semantic zoom levels, provided by different views.

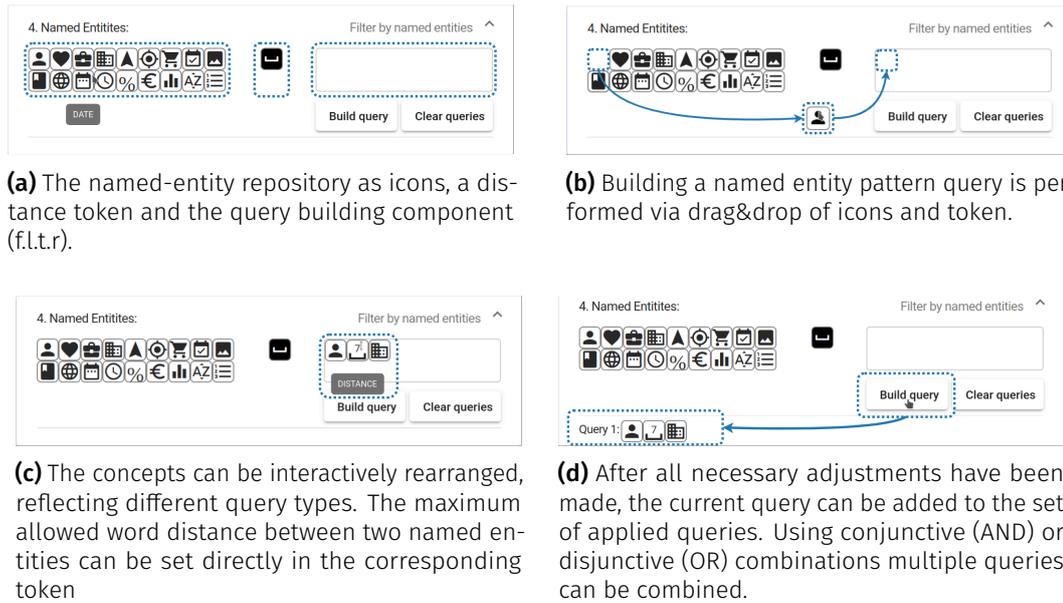


Figure 7.5: Illustration for the step-by-step process for creating a named entity relation pattern query, searching for two concepts that occur within a given word length of each other.

7.5.2 Level Properties

Each level can have its own properties, accessible through a property pane  on the right in the main user interface. The standard task levels offer controls like cutoff values, include/exclude lists, or time sliders. The dynamics level offers restrictions on the individual properties of conversational dynamics [SFS⁺19]. Here, we want to describe one more complex property that can be offered on the property pane  for custom levels, using the **Thematic Level** as an example: a visual query interface for thematic searches using named entities. To generate the named entities, we employ a pre-trained model from spaCy [Hon19], containing a set of 18 named entity categories. The interface itself is shown in Figure 7.5, illustrating the individual components and the step-by-step process for creating a sample named entity relation pattern query, searching for two concepts that occur within a specified word distance.

As shown in Figure 7.5a, our visual query interface consists of three main components: A repository containing a set named entities, such as persons, appointments, or organizations, a special token to allow distances between naming entities, and finally, a query building component. To build a named entity pattern query, a user can drag one or multiple individual named entities from the repository and (optionally) the special token into the query building component, as highlighted in Figure 7.5b. The concepts can also be rearranged inside the query building component, reflecting different query types, like single concepts, a chronology order of concepts, or

distances between the concepts. For the latter, the maximum allowed word distance between two named entities can be set directly in the corresponding token in the query builder. In the example shown in Figure 7.5c, the maximum allowed distance between named entities is set to seven words. After all necessary corrections and adjustments have been made, the current query can be added to the set of applied queries (Figure 7.5d). The query shown only serves as a single example of a named entity query, where other types are possible.

7.5.3 Machine-Learning-based Retrieval System

The design, using multiple complementing levels, allows for cross-level search, with level-specific properties that can act as filters, acting together to form a global filter. This already allows for more powerful and interrelated search tasks than the application of individual levels alone. However, depending on the types of filtering defined and their interactions, the possible settings might overwhelm domain experts. Therefore, we propose that each level can output a feature vector for a communication event. Level-specific vectors can be combined to a single, large feature vector used for classification purposes in a user-steerable machine-learning level. Although progress has been made to use deep learning efficiently by reducing training time [SK16] and improve explainability [GMT⁺18], their usage is still problematic when requiring (theoretical) traceability, for example, due to legal constraints. Consequently, as proposed in the literature [MQB19], we employ a rule-based approach based on a random forest model. However, while this can fulfill legal requirements, from a perspective of lay use, a random forest's decisions might still be tough to understand. It has, however, the additional benefit that the training size can be relatively small (usually much less than a few dozens), making it suitable for an interactive application, while the examples can easily be collected by the users themselves, based on their expert knowledge. A user can train individual automatic classifiers that support him on specific tasks and modularly combine them to compute overall predictions.

The selection of training examples for the classifier happens interactively. A user can label communication in a binary way as relevant or irrelevant to perform a semi-automatic classification of communication into user-defined classes. An example of such a selection is shown in Figure 7.6. Such a trained classifier can then perform the binary classification for all other communications, acting as an additional high-level filter. Since we use a Random Forest Classifier, we can model the uncertainty for the prediction, which is useful for thresholding. This also allows presenting ambiguous communication, where the classifier is very uncertain, to the user for re-labeling, allowing for an interactive optimization. To separate between this semi-automatic retrieval system and manual level property settings, communication that is filtered

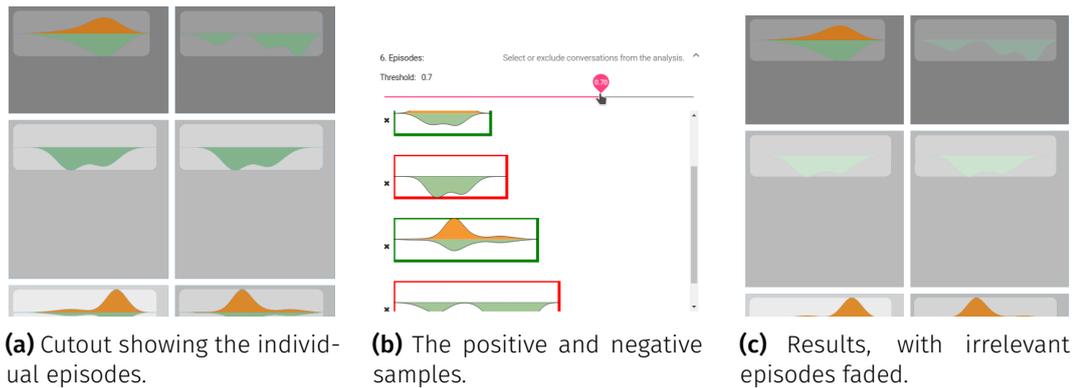


Figure 7.6: By providing feedback, users train ML models to identify relevant conversational dynamics in episodes. Here, the aim is to identify episodes in which the selected groups start the conversation, leading to a discussion of both entities.

Provenance History:

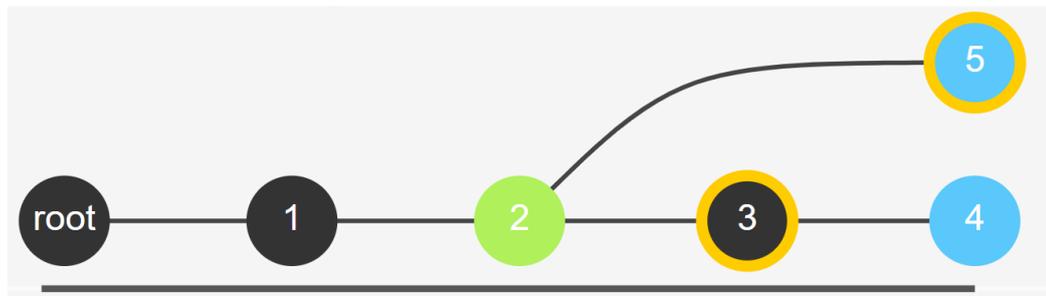


Figure 7.7: Provenance history component showing previous filter states as nodes in a directed acyclic graph (DAG). The **currently selected node** and the **leaf nodes** are specially colored and important states can be **starred** for easier distinction. The states allow for more flexible navigation, to revert from dead ends or branch off to a different analysis direction.

out by the automatic system is only faded out based on a variable threshold, but not hidden completely.

7.5.4 Analytical Provenance

To support experts even in lengthy and complicated investigations, we offer a provenance history component, which is displayed in Figure 7.7. Since explainability is relevant, for example, in court cases, decision making and a record of the obtained results must be preserved. The provenance history contains previous filter states as nodes in a directed acyclic graph (DAG) to allow for more flexible navigation compared to a linear history [HSS⁺]. Important states can be visually starred. The user can navigate between different states, go back to previous results, or branch off as a new starting point for further analyses. This enables the user to continuously verify and retrace results, which is especially advantageous for creating trust in the

user's results. Further, the individual filter states can be bundled into a physical report and thus archived. In this way, the analyst's results can be reproduced, retraced, and explained, even after an extended period and/or to third parties. Since each of the steps and all of its belonging information and metadata and the obtained results can be reviewed and analyzed independently, allowing for explainability of the results obtained.

7.6 Evaluation

To demonstrate the effectiveness and improvements compared to existing approaches of the visual exploration of communication behavior in COMMaid, we conduct an expert assessment of the prototype while additionally conducting a case study together with a small-scale user study later on. As communication data, we use the largest publicly available source, the Enron dataset [KY04], encompassing 517 431 messages from 151 users.

7.6.1 Formative Domain Expert Assessment

The assessments were conducted by demonstrating the prototype to six domain experts (LEA 1–3, RS 1, SI 1–2).

Expertise LEA 1 is a criminal investigator at a European law enforcement agency with extensive experience in the field, including communication and network analysis with graph-based visualization using systems like IBM i2 Analyst's Notebook, Pajek, and Gephi. LEA 2 is a criminal investigator with no prior experience using graph-based visualizations. He focuses on the communication content, using forensic tools like Cellebrite and IBM i2 Analyst's Notebook, which is laborious. LEA 3 is a senior judicial commissioner in law enforcement with extensive experience in digital investigation techniques. He is aware of the graph-based approaches used within his unit, but has limited own expertise. RS 1 is the head of a university-affiliated institute for policy and security research, a full professor and senior researcher working on government projects. He has worked with graph-based visualizations and approaches for over 20 years, for instance, for bibliometric investigations. SI 1 is a senior project lead within the security industry for developing investigative solutions for LEAs. SI 2 is a junior solutions specialist within the security industry and worked on visualization techniques, including graph-based visualizations, for criminal investigations before. All but one (SI 2) of the experts have more than 15 years of experience in digital and criminal investigations.

Methodology The expert assessment was conducted as a formative evaluation taking 90 minutes, with the experts observing, commenting, and asking questions on the system, while provided with an user sheet. They were also allowed to steer the

exploration by requesting specific actions. As such, a formative evaluation without direct usage usually cannot replace the benefits of a full user study; therefore, we additionally describe a case study in Section 7.6.2 which we used to conduct a user study with two further domain experts later on, as described in Section 7.6.3. The experts were first given a ten-minute introduction about the aim, which is aligned with the tasks defined in Section 7.1: facilitating the visual analysis and exploration of large amounts of communication data using different visualization and filtering methodologies simultaneously to structure and reduce the search space to a manageable size for inspection by an analyst. The prototype was then showcased and explained during 30 minutes of interactive demonstrations, where the experts actively asked questions and commented on the presented aspects. They include using the network overview to detect promising connections and explore individual details as well as communication episodes in level-specific visualization through drill-down. Further, we presented and then debated the different interaction techniques, then the available filters to reduce the selection, and finally the example-based machine learning retrieval classifier.

The interactive session was followed by a structured interview, taking about 50 minutes, using a set of 29 prepared questions about various aspects of the approach. This interview was interluded by interactively presenting aspects in the demonstrator when requested by the experts. The session aimed to elicit experts' opinions on the system's design and interaction decisions and identify aspects they find helpful or prone to misinterpretation. Further, we were interested in how they would apply these methods in their specific workflows and criminal investigations in general. The findings of these assessments are described in the following.

Findings All the experts state that both the approach of using a **matrix-based overview visualization** and using a semantic zoom for more details is a new approach in their domain. For example, according to LEA 1, he has "always worked with graph tools" so far and thinks of our technique as "very interesting and helpful". All experts think that a matrix-based visualization is superior to graph-based approaches in "terms of scalability" (SI 2) and displaying "supporting information" (RS 1). However, both LEA 1 and LEA 2 recommended the matrix columns to be freely reorderable.

Regarding the **semantic zooming**, the experts are familiar with such a concept from everyday applications like digital maps. For communication analysis, some did not expect this functionality at first (cf. LEA 2). However, it supports their work and is an excellent way to drill down to "go into the raw data" (cf. LEA 2). RS 1 judged the semantic zoom as intuitive, but expected that - instead of the communication structure and content - more information about the "importance of relations" (RS 1) is shown. RS 1 proposes to include such information as another level, as the design

is “flexible enough” (RS 1). For example, an **analysis levels** for centrality analysis could be added.

In terms of **filtering**, the experts are happy to have the standard task functionality included. However, more advanced concepts like the semantic named entity search were “unexpected” (LEA 1, LEA 3). In their previous experience, the LEA’s were only able to search using lists of keywords and were “never able to search for concepts” (LEA 1). They regard this functionality to have much potential, as it allows for “more generic” (LEA 1) and high-level search terms. As the prototype’s current implementation restricts them to AND-based queries, all experts state they would like to combine the queries more freely using Boolean logic.

Regarding the **machine-learning based retrieval** of communication episodes, all experts agreed that detecting related and sequential communication is “important for contextual information” (SI 1). The visualization as density-based communication amount can be intuitively understood by all experts. LEA 2 regarded it as beneficial that the detailed raw data from a communication episode can be inspected within the visualization. Going on from there, the ability to train a machine learning model by giving communication episodes as examples is viewed as “opening up new possibilities” (RS 1). Domain experts especially favored that **arbitrary features** can, in principle, be used for the machine learning model (extending upon those we defined above), which can be communication “based on text, audio, pictures, geographic information systems, or combined with graphs” (RS 1), using additional levels. Therefore, the features itself “do not matter much, as the user has to define them” (RS 1), making the flexibility of our approach “very broad” (RS 1).

With such broad applicability, the explainability and retracing of the results are “an important issue. If an analyst has a result, he needs to explain how he ended up there” (LEA 3). This explanation is simplified tremendously “with the generation of a step-by-step report” (LEA 3), as producible from our **provenance** history, whereas currently, “analysts have to write detailed accounts on how they got to the information” (LEA 3) and justify it each time in writing.

In terms of practical usage, they had “no ideas for additional [conceptual] features” (LEA 1), except for the inclusion of centrality measures. But the potential possibilities with the framework are almost “overwhelming at first” (SI 1). The **applicability** of the presented approach is not restricted to a narrow use case as the one presented. Therefore, the system is “broadly applicable to multiple domains where you have bigger groups of communication data. Data that [law enforcement experts] often have to deal with. For example, organized crime, financial crime, or terrorism” (LEA 3). All the other experts share this view. Indeed, the presented system is “a beginning of an interaction platform where you can combine other logics as well and [which] offers many possibilities” (RS 1), providing custom analysis levels specific for your needs (cf. RS 1).

7.6.2 Case Study

During the formative expert assessment, the experts did not interact on their own with the system, as we were only able to secure a limited, non-individual amount of time with them. To compensate for bias, we additionally conducted a small-scale user study with two further domain experts (LEA 4 and RS 2). For this, we describe a possible case study, highlighting the benefits of an holistic approach.

Methodology and Case Study The case study is based on an artificially financial fraud use-case to identify **senders**→ and **receivers**← of relevant communications. The aim is to discover those persons 👤 which, during the **first nine months of 2001** 📅, disseminated knowledge about **legal** 📄 issues involving **persons** 👤 in combination with **organizations** 🏢 in **California**, and then identify the only person 👤 who received information from all of them. The experts were given a system manual and a short written description of this task. Such a task often occurs in real cases, but is not well supported by existing approaches. Instead, one often relies on keyword list using domain-specific ontologies, requiring manual work to create and search through the results. Here, we present how an **holistic analysis** of **network structure and dynamics** 🌐, **metadata** 🔑, **keyword-based search** 🔑, and **semantic concepts** 📖 can address this task.

The expected solution was to map the task conditions to the analysis levels and their properties: namely **specific time ranges**, applicable **field offices**, and concepts that are relevant for the search (**persons** 👤 🗃 **organizations** 🏢, and **legal topics** 📄). In a second step, when they have identified the persons 👤 that disseminated the knowledge, they were expected to make the mental transfer to restrict the view to those participants as **sending persons**→ and examine which participant is the only **receiver**←. The successful completion of the task was measured by checking the name of the identified users.

7.6.3 User Evaluation

The case study in the previous Section was used to evaluate the system as part of a user evaluation. During the task, the users were undisturbed to explore and try out the system freely. They themselves could decide how long they want to train and check out the system, before actually starting. It was the first time they saw and used the system for both domain experts, having no prior experience with it and having received no introduction, so we could do a blind test. After completion, we interviewed them shortly about their experience. The participants of the user evaluation are described in the following, before we discuss the findings.

Expertise LEA 4 is the head of the big data department of a federal governmental agency supporting criminal investigations. She advises law enforcement agents on the applications of artificial intelligence to criminal investigations. RS 2 is a

senior scientist at a federal government research institute with over ten years of experience in communication analysis and terrorist investigations.

Findings Both experts were able to successfully complete the task within 15 minutes. The interface and **interaction concepts** were described as “intuitive [and] self-explanatory” (LEA 4). The first expert was surprised to be offered a search for concepts and initially tried to use the keyword-based search level for conceptual searches instead of the visual query interface, while the other did not make this mistake. We conclude that some users would benefit from a more hands-on explanation. LEA 4 took five minutes longer to complete, exploring the options and results in between, but ultimately solved it like RS 2 (see solution above), who went ahead directly. The system offers “helpful” (RS 2) drill-downs of visualizations and is intuitive and straightforward (cf. RS 2) to use. Both domain experts noted that the system naturally supports the investigative **workflows**, and the interaction design combined with the documentation is sufficient for working productively and getting relevant results. Compared to their existing systems and workflows, the system provides a significant benefit in analytical capabilities. Most notably, it allows the **simultaneous application** of different search methodologies to support cross-matches. This allows for more powerful queries, in contrast to manually merging separate results.

7.7 Discussion and Future Work

During the formative expert assessment and the study, we received several proposals on how to extend our approach further. Leaving out expected requests for a research prototype like more supported data imports, we instead focus on the core functionality of the approach by discussing the limitations and the broader applicability as well as the context of future work. For our prototype approach, we adapted the generic blueprint of modularized communication analysis to the case study by providing two example levels.

Of course, the system can be extended modularly with **further analysis levels**, for example, those featuring graph centrality measures, community detection, leveraging specific meta-data like locations, or more specific content analysis modules based on linguistics. Currently, this requires modification in the source code, but a plug-in architecture would be conceivable. Regarding the overall **matrix visualization**, the challenges and limits of such an approach for communication analysis have partly been discussed in previous work [FAS⁺20]. When tasks are primarily focused on network structure analysis, a classical node-link-diagram-based approach might be more suited. Alternatively, one could consider a coordinated view approach and adding a synced node-link-diagram. As space is limited, making both matrix and node-link components user resize-able would allow for a task-adaptable interface.

Intelligent layouting of the node-link diagram based on different parameters (centrality, connectedness, meta-data) would support visually finding patterns, which can be further analyzed in the matrix view and vice versa.

A challenging step for future work is to investigate how this approach can be used for the analysis of **multi-party conversations**. So far we duplicate messages with multiple recipients, which partly destroys the group aspect. Future work could investigate how hypergraphs allow to capture such scenarios or how additional zoom levels could be leveraged for group communication analysis. Also, we assumed some practical restrictions to describe the network by requiring receiver, sender and timestamp information for each message. As some analysis steps require those, incomplete entries (like unknown recipients) cannot be represented. Another aspect of the prototype is the **visual query language**, where it could be valuable to extend the grammar and support nested queries visually.

Finally, the expert interviews and the user study resulted in a positive response to our workflow and prototype system. However, we are well aware that the sample group for our formative experts' assessment and the user **study was quite limited**. To achieve statistically more accurate results and broaden the perspective, the study could be extended with more participants as part of future work. This explainable and accountable decision making is not only relevant in the security domain, in which the case study and expert assessment were conducted. Indeed, the experts think the applications are not limited to such a narrow set of criminal communication investigations but can be applied to communication data in **other domains**. One different application would be in the business intelligence domain. The system could be applied as a search and retrieval mechanism to search for hidden, decentralized knowledge contained in business documents and communications. This knowledge can then be linked and extracted into centralized knowledge management systems, allowing for more efficient management structures and avoiding redundancies, making the processes more accountable.

7.8 Conclusion

So far, most interactive, automated communication analysis approaches focus either on the network aspects or on the content, in contradiction to communication research. As such, the individual or isolated analysis may not suffice to capture the full available information and may lead to less effective, incomplete, and biased results. Further, it can increase the struggle experts face when articulating their domain knowledge, not leveraging their full potential.

We address this challenge by arguing for and discussing a holistic approach to communication analysis, simultaneously applying both methods, allowing for more structured and detailed analytical capabilities. To help domain experts deal with

the complexity of modern communication data, we present COMMaid, a blueprint for a visual analytics-based communication analysis system that offers a wider approach, providing a tight coupling between the network and the content analysis aspects, building on individual levels and supported by a machine learning-based retrieval system.

We provide two extendable levels as an example for network and content analysis each, covering dynamics modeling (based and extend from our previous work [SFS⁺19]) and semantic text analysis. We leverage ideas from hypergraph analysis [FAS⁺20] for a multi-level matrix-based visualization design to integrate those levels in a single interface. However, we specifically tailor and adapt this idea to communication analysis by providing specific visualization levels to support domain experts in their mental understanding during exploration and allow them to answer more detailed questions about communication behavior and structure, including identifying individual communication episodes. Combining network and content aspects in a single visualization allows for maintaining overview and focus while eliminating demanding context switches, rapidly exploring large search spaces, and providing details on demand. The realized techniques allow the simultaneous analysis of network and content aspects, like properties, conversational dynamics, or conceptual content, to refine the search and supports cross matches.

We evaluate our approach in one case study and through assessments with law enforcement experts using real-world communication data. The results demonstrate that our system surpasses existing solutions, enabling the effective analysis of large amounts of information in a targeted and integrated way. The experts regard this approach as a novel and promising way for a more meaningful communication analysis that can readily be applied to comprehensive analytical tasks as encountered in practical applications. While we focused on communication analysis for law enforcement as driving application, many tasks in communication analysis are similar and, therefore, our methods are more generically applicable to a wider variety of domains, like digital humanities or business intelligence. By bridging this gap between network and content analysis in semi-automated communication analysis systems, we aim to pave the path for a more holistic approach to communication analysis.

δηλον γάρ ὅτι τῶν μερῶν ὄντων οὐδέν κωλύει τὸ ὅλον μὴ εἶναι, ὥστ' οὐ
ταύτῳ τὰ μέρη τῷ ὅλῳ
For there, clearly, you may have the parts and yet not have the whole, so
that parts and whole cannot be the same.

— Aristotle, *Philosopher*

8

MULTI-CASE: A Multimodal Visual Analytics Technique

In the previous chapter, we presented COMMAID as a blueprint for a holistic analysis of communication data. In the following, we want to extend the general idea for a contextual, nuanced analysis by going further and not only describing the (textual) analysis of communication data but embedding the analytical capabilities in a broader concept: Primarily challenges faced by current approaches are related to the consideration of related multimodal data, the transparent inclusion of AI models, and the implications as well as difficulties in actually designing such systems from an ethical and privacy perspective. AI-driven models are increasingly deployed in operational analytics solutions, for instance, in investigative journalism or the intelligence community. We discussed the ethical and privacy concerns as well as the challenges involved previously in [Chapter 3](#) and will apply the lessons learned in the design of the following technique. Further, we have highlighted the need for a holistic analysis both in [Chapter 2](#) and [Chapter 7](#), and in the following, we describe a technique to efficiently combine heterogeneous data sources for both holistic as well as multimodal analytics, which also includes related data like audio, video, images or other data types. To tackle the challenge of multimodal analytics, we present MULTI-CASE, a holistic visual analytics framework tailored towards ethics-aware and multimodal intelligence exploration, designed in collaboration with domain experts. It leverages

Contents

8.1 Multimodal Communication Analysis	172
8.2 Related Work	176
8.3 Methodology: Model Development	178
8.4 System Design	181
8.5 Evaluation	186
8.6 Discussion and Future Work	198
8.7 Conclusion	201

an equal joint agency between human and AI to explore and assess heterogeneous information spaces, checking and balancing automation through Visual Analytics. MULTI-CASE operates on a fully-integrated data model and features type-specific analysis with multiple linked components, including a combined search, annotated text view, and graph-based analysis. Parts of the underlying entity detection are based on a RoBERTa-based language model, which we tailored towards user requirements through fine-tuning and published as open-source. An overarching knowledge exploration graph combines all information streams, provides in-situ explanations, transparent source attribution, and facilitates effective exploration. To assess our approach, we conducted a comprehensive set of evaluations: We benchmarked the underlying language model on relevant Named Entity Recognition (NER) tasks, achieving state-of-the-art performance. The demonstrator was assessed according to intelligence capability assessments, while the methodology was evaluated according to ethics design guidelines. As a case study, we present our framework in an investigative journalism setting, supporting war crime investigations. Finally, we conduct a formative user evaluation with domain experts in law enforcement. Our evaluations confirm that our framework facilitates human agency and steering in security-sensitive, AI-supported analysis processes while addressing ethical and privacy concerns and providing much-needed analytical capabilities.

This chapter is based on the publication [FMJ⁺24] and major parts of the following sections have appeared in:

- [FMJ⁺24]: **Maximilian T. Fischer**, Yannick Metz, Lucas Joos, Matthias Miller, and Daniel A. Keim. “MULTI-CASE: A Transformer-based Ethics-aware Multimodal Investigative Intelligence Framework”. In: (2024), pp. 1–16. DOI: [10.48550/arXiv.2401.01955](https://doi.org/10.48550/arXiv.2401.01955).

For a statement of the scientific contributions, as well as the division of responsibilities and work in this publication, please refer to Chapters 1.2 (p. 22ff) and 1.3 (p. 26ff), respectively.

8.1 Multimodal Communication Analysis

AI-driven models have gained wide popularity over the last few years and have been applied successfully in numerous fields, such as natural language processing (NLP), computer vision, or predictive analytics. Given this general trend, AI

models are increasingly needed [Mit21] and deployed in operational intelligence solutions [BYW20; Gan21]. Corresponding application domains, such as investigative journalism [BDG⁺19; Str19] or the intelligence domain [BYW20; HM21; OS23; M⁺g23], are particularly interesting due to their unique set of distinct challenges. Intelligence analysts often face the task of combining numerous, heterogeneous pieces of intelligence, often tainted with uncertainty and conflicting information, forming an incomplete picture. As discussed in Chapter 3, the first set of challenges in this regard is related to **ethical** [Shn20] and **privacy concerns** [Mit21] due to the sensitive nature of the data and operations involved [Rig19] and the high stakes in case of errors [Asa19; ADP⁺22]. Simultaneously, these domains offer opportunities for increasingly automated, tailored systems to deal with incomplete and tainted information. This is particularly the case for **heterogeneous** and **multimodal analytics**, a second area in which existing systems often lack in functionality [FSS⁺21; FDS⁺22b].

The analysis of **individual modalities** in isolation—like network structure of the participants, named entity detection on the content, or time series analysis of the individual message intervals—often comes with limited views on the underlying information with consequences for the derived intelligence. Not considering these aspects can reduce trust in AI systems, favor prejudices and mistakes, and also lead to legal consequences. Further, isolated analysis requires human knowledge and intervention to semi-manually find hidden cross-matches between the modalities—a task where computational support can be highly effective, reduce domain discontinuities, and place less additional workload on the users [FSS⁺21]. This becomes even more important when users are no machine learning experts, thus sometimes having unrealistic expectations or misplaced trust in the systems [FHJ⁺22; Mit21]. This can be the case for (business) intelligence analysts or investigative journalists, after which we modeled a case study (see Section 8.5.1).

This study is based on widespread **tasks** in intelligence, identified by the UNODC [UNO11], which aims to answer the typical six questions: *Who? What? How? Where? Why? When?* Based on these six questions, the UNODC authors identify three common analysis tasks and methods that typically enable the answering of these questions in relevant investigations: (1) *link analysis*: searching and identifying relationships between specific entities such as persons or organizations, but also objects, locations, or events, (2) *event analysis*: correlating actions or locations alongside their timeline order, (3) *flow analysis*: understanding the connectedness as well as cause and result, for example, the flow of commodities (geolocation for physical goods or transfers of money) or the propagation of knowledge. Other tasks described in the report involve the identification of activities, frequencies, or general data correlations. These tasks can be primarily achieved through four main methods: (a) keyword and semantic-based searches on text or transcripts to understand the context or find entities, (b) (social) network analysis to find connections and

relations, (c) meta-data-filters to restrict, for example, locations, and (d) time-series analysis, for example, to identify particular communication patterns. However, these modalities should not be considered to work in isolation but contribute individual perspectives for corroborating, enhancing, and setting each other in context. For example, to attribute war crimes in our case study (see Section 8.5.1), our journalist Alisa leverages semantic analysis, geolocation, link-analysis, and time-correlation together with several other methods to achieve her objectives.

Our **objective** is to tackle the existing shortcomings in ethical and multimodal analysis for intelligence by presenting a framework for holistic communication analytics. Many specific solutions have been proposed, but the integration and combination have received less attention. In [Chapter 2](#) and [Chapter 3](#), we have detailed the data and problems faced in intelligence analytics: the need for heterogeneous data analytics capability due to the diverse set of intelligence received. The different data types and scenario stakeholder groups like data subjects, software providers, civil society, and governmental authorities with their different branches with all their conflicting interests. Their requirements and tasks, which we also revisit below, as well as the benefits and possible designs of visual analytics applications.

Our contribution is not intended as a fully-fledged analytics system but as an exemplary *framework* for a holistic, multimodal approach to intelligence and its assessment. Therefore, we dedicate significant time towards a comprehensive *evaluation* (see Section 8.5), encompassing multiple perspectives, i.e., ethical aspects, capabilities, and practical considerations through use cases and expert studies.

Based on lessons learned in [Chapter 2](#), [Chapter 3](#), and [Chapter 7](#), we aim to enhance the analytical capabilities in semi-automated digital intelligence analysis, making the following **contributions**:

Contributions

- MULTI-CASE, an *integrated visual exploration framework* (see Fig. 8.1) tailored towards ethics-aware *multimodal* intelligence analytics in investigative journalism or criminal investigations.
- A RoBERTa-based NER **transformer model**, derived by fine-tuning on GottBERT [STT⁺20] alongside intelligence-specific training data, which we both open-sourced at osf.io/eap4r.
- An extensive **case study** showcasing MULTI-CASE in the context of *war crime investigations* together with a **classification assessment** of its **capabilities** [FDS⁺22b] and **ethics design** [FHJ⁺22].
- A formative **expert evaluation** with eleven domain experts in different law-enforcement areas, validating the approach's advantages and highlighting areas for further improvement.

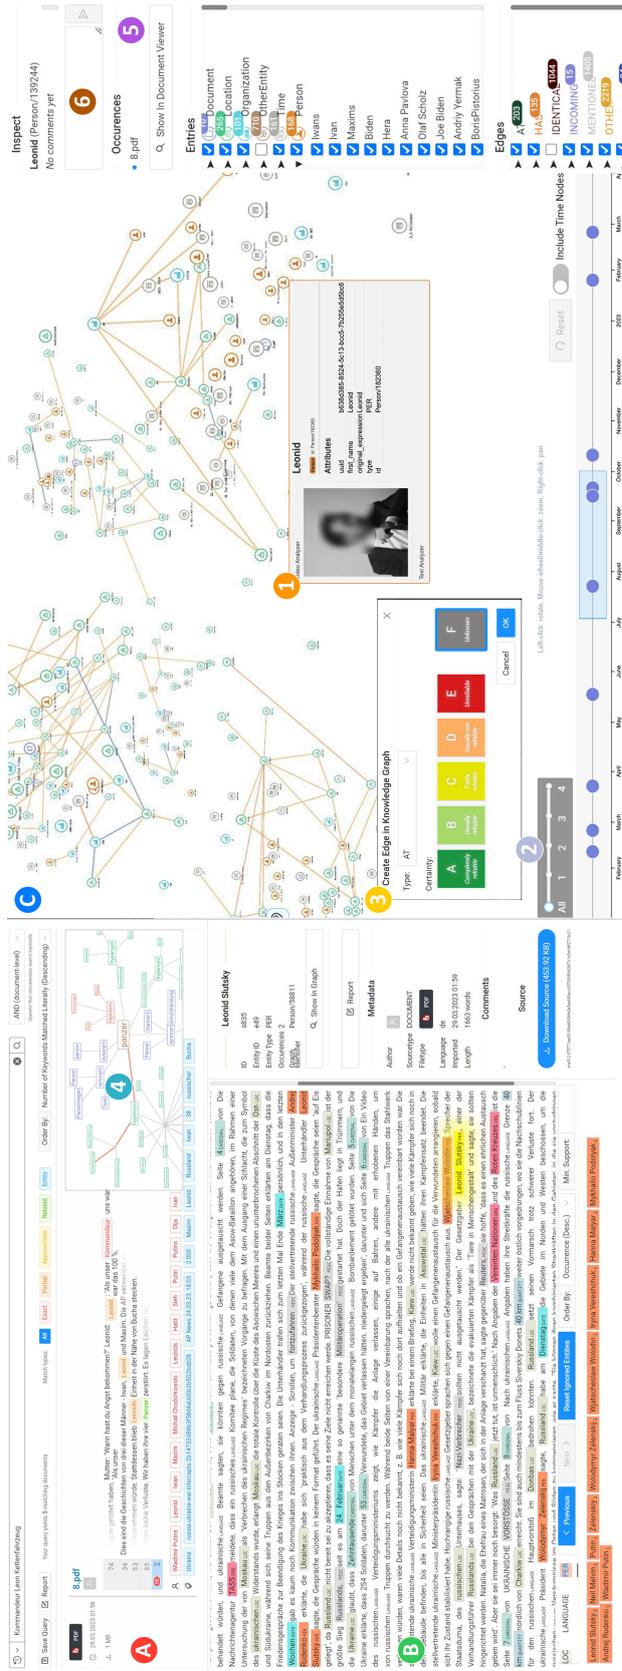


Figure 8.1: MULTI-CASE: A holistic visual analytics framework tailored towards ethics-aware and multimodal intelligence exploration. Built upon a fully-integrated data model, it features type-specific, graph-based analysis through individual models with multiple, linked components: **A** a combined ontological search and result interface, **B** an interactive textual view, and a **C** knowledge graph interface. Additional components (not shown) include more specialized modules like video or audio analysis. The interface facilitates **1** in-situ contextualization across modalities, **2** graph neighborhood explorations, **3** relevance scoring for accountability and oversight, **4** transparent source explanations, **5** integrated navigation, and **6** collaborative user participation.

With this contribution, we fill a gap in bringing state-of-the-art performance to applications by providing an explainable visual exploration framework for multimodal intelligence analytics. We consider our contribution primarily in the combination of existing visualization and visual analytics methodologies suitable for this domain and their detailed assessment in the context of the unique challenges faced. Thereby, we aim to provide more insights into the often opaque workings in the intelligence domain, furthering research and a critical discussion.

8.2 Related Work

The research on **multimodal visual intelligence** analysis is sparse. While there is significant literature on intelligence analysis in general [UNO11; GWA⁺17] and some requirement studies for general intelligence analytics tools exists [EPT⁺05; Sch05; FDS⁺22b], actual tool descriptions are rare. If a paper evaluates an actual approach, their findings primarily focus on user acceptance while ignoring capabilities or interactive visualizations since tools are often classified and not even named publicly [Dha17].

Research on some of the underlying techniques itself, for example, classical **Named Entity Recognition** (NER) as the foundation for comprehensive tasks like entity linking, is much more common. Techniques evolved over time from using rule-based to more statistical systems [NS07]. Traditionally, NER relied on annotated corpora, which posed challenges for domain transfer and new label tasks, with brittle results [NS07; LSHL20]. However, with the advent of deep learning-based approaches, such as BERT [JMCK19], the landscape has changed, and transfer learning (i.e., adapting pre-trained models to shorten training times for new tasks) can cope with much smaller amounts of annotated text. This has significantly improved the adaptability and efficiency across various domains and tasks, making knowledge transfer and few-shot labeling easier [JMCK19; LDS17], which can be leveraged in investigative tools.

Similarly, advances in **ethical design** [BM19; FHJ⁺22], like the concept of providing guidance [SCE22], visualizing hidden uncertainties [ZLVV22], or ensuring provenance [Cor19] as well as **privacy considerations** [FHJ⁺22], like selective masking [THW⁺21], federated learning [GKN17], or data perturbation [SMC20] have been made. Also, **insular solutions** like Pajek [BM02] for social network analysis, Maltego [SC21] or InSight2 [KKG20] for link analysis, or Cosmos [DWM⁺19] for semantic text analysis exist but do not combine modalities.

Within the visualization community, **multimodal multimedia analysis** [ZW14] can be considered partly related: Several approaches have been proposed to consider different aspects of multimedia content simultaneously, like the presentation styles and techniques [WQ20], the emotional coherence [ZWW⁺20], or the automation

of explicit content through video moderation [TWW⁺22]. While these approaches propose valuable insights into how (primarily visual) media can be analyzed and set into context, many of the approaches target very specific applications, and very few in this domain truly support a holistic approach to analyzing *generic* pieces of intelligence, which also includes text-based information. Further, Zahalka and Worrying presented a pathway to comprehensive multimedia analytics, detailing a general four-tiered multimedia analytics model and discussing it alongside how it may support addressing the semantic and pragmatic gap encountered in actual systems [ZW14]. This follows a similar overall direction as our research, however, with one particular difference: The model is applicable in general for the analysis of multimedia data and also with a particular focus on such data, for example multimedia collections of images. While some aspects overlap, these collections of images do not necessarily have a underlying storyline, may come from any collection mechanism (e.g., underwater camera), and the model primarily focuses on a multimodal analysis of multimedia with additional metadata (e.g., annotated text or features). Our approach instead focuses primarily on communication between humans, emphasizing much more the interactive aspects of the information exchange via various modalities over time.

The research on leveraging **visual analytics for intelligence applications** [DGLR09; KGS09; KS11; LKT⁺14] had its prime in the mid-to-late 2000s, with frameworks such as VIM [KC04] or Jigsaw [SGLS07]. Both primarily focus on text documents (and not so much multimedia), and only a few approaches [FSS⁺21] were proposed later on. Therefore, this area seems to be one of those few domains where commercial research has outpaced academic, scientific research for now.

In the context of actual usage—also for commercial systems—we surveyed related communication analysis systems [FDS⁺22b], where we identified four publicly known intelligence systems in wider use: DataWalk [Dat20] and Nuix Discover / Investigate [Nui20] are sometimes used, while the market leaders are IBM i2 Analyst’s Notebook [IBM20] along with Palantir Gotham / Foundry / Meta-Constellation [Pal20]. While they cater to government applications, parts are commercially available and are used by international banks, advertisers, manufacturers, telecommunication providers, media organizations, and NGOs [Pal20].

To our knowledge, no new visual analytics approaches to intelligence have been publicly proposed since our recent survey of AI-driven intelligence applications [FDS⁺22b], also available as an **interactive browser** at <https://communication-analysis.dbvis.de>. Regarding practical usage, the ongoing shift from IBM i2 to Palantir seems to accelerate. Palantir’s solutions (in particular Meta-Constellation) are also employed effectively [Sco22] by Ukraine in its defense against Russia in coordinating their military.

The **academic research** on this topic has been falling short, with problematic consequences for accountability and oversight, which has also been realized by some key stakeholders. For example, in the European Unions Horizon 2020 funding period alone, projects such as ASGARD (700381), MAGNETO (786629), STARLIGHT (101021797), COPKIT (786687), and AIDA (883596) (some still ongoing) have been funded, although preliminary results show insular capabilities. For the upcoming Horizon Europe funding period, several calls have been proposed (e.g., HORIZON-CL3-2023-FCT-01). With slight deviations, they all aim to increase analytical big data capabilities for law enforcement. In the US, similar research is often conducted by national laboratories but mostly remains classified.

While many visualization approaches can be leveraged for intelligence, only few consider the combination of challenges faced in this particular domain, including the inherent uncertainty and inter-modality, while even fewer evaluate them consistently and publish the results, which is the goal of this work.

8.3 Methodology: Model Development

In [Chapter 7](#), we have presented a matrix-based, holistic communication analysis framework through semantic zooming. As our studies have shown, however, despite the potential benefits in scalability, matrices are uncommon for many analysts, which are used to graph- and relationship-based visualizations. Further, semantic zooming is space-limited in the amount of context information in the upper layers. We, therefore, aim to explore an *orthogonal design*, with *two key advancements*: (a) Following a similar modular approach, we leverage a more powerful **fully-integrated data model** (structuring and relating the intelligence information pieces) that also supports multimodality. (b) Instead of matrix-based semantic zooming, we use a **graph-based overview** with several linked views and integrated specialized views.

This decision is based on the task descriptions and requirements described in the UNODC report [UNO11] described above, as well as feedback from several domain experts in law enforcement, which state the following three **user requirements** for such a framework: (1) A centralized, multimodal platform for collaborative case working. (2) Assistance in labor-intensive tasks such as big data analytics. (3) Transparency and reliability.

This reflects their need to *work collaboratively* on a case together with their colleagues on larger investigations, needing to share results or to leverage knowledge generated by colleagues investigating specific aspects of a case by collaboratively working on a shared data space and being able to access the information in-situ. Due to the sheer volume and sometimes repetitive tasks, support by automation and AI is considered essential while being reliable and understandable. All the while, the analysis steps taken need to be transparent and reproducible for accountability.

Guided by these overall principles, we further justify individual design decisions and capabilities while describing the system design in Section 8.4.

One central aspect of intelligence analytics is the analysis of communication [FSS⁺21]. However, common international NER labeling schemes (e.g., PER, ORG, LOC, OTH) often do not meet the specialized requirements for investigations since they are too ambiguous and not specialized enough, requiring more narrow tag categories [MWBB13; RR09]. In practice, specialized **NER model development** for semantic understanding is still challenging, with many pitfalls, although the attention-based transformer architecture [VSP⁺17] has significantly increased the accuracy compared to previous neural models. Therefore, as part of this work, we track the necessary steps for training and deploying transformer models, including interactive tools for labeling, while also highlighting major lessons learned. The necessary steps range from choosing a suitable base model, preparing representative training data, then training and evaluating, to finally supervising and validating the model in deployment and adapting it in the face of changing language patterns, terms, or requirements. As a result, we provide a strong baseline NER transformer model with a large set of relevant entity labels to simplify future applications. For the underlying language model, we considered existing models from the Huggingface transformers [TLV⁺19] library based on evaluation performance on the GermEval14 dataset [BBKP14], a well-known dataset for German NER recognition. For German natural language processing, we considered two language models: The RoBERTa-based GottBERT [STT⁺20] and *BERT-base-german-cased* [Sch22] based on the original BERT transformer architecture [JMKK19]. Additionally, we chose a strong multi-lingual baseline (XLM-RoBERTa) [CKG⁺20].

In general, the creation of specific training datasets, for example, through **labeling** of domain-specific datasets, is often tedious and error-prone. Therefore, we implemented an interactive labeling tool that is compatible with the MULTI-CASE framework, allowing us to label and subsequently review a given document collection on a large scale, facilitating the easy creation of ground truth training datasets in specific domains, like intelligence. This is particularly relevant in our application because it utilizes a large set of custom-named entity labels for domain-specific analysis. Many non-English models only provide standard categories like *PERSON*, *LOCATION*, *ORGANIZATION*, and *MISC*. However, based on expert feedback, custom categories like *EVENT* or *PRODUCT* and more fine-grained time and numeric labels were introduced, with the full list shown in Table 8.1. We provide an enhanced, re-tagged version of GermanNER alongside our model at osf.io/eap4r.

For the **training**, we apply a train/validation/test split of 70/15/15 of the full mixed dataset (domain-specific and re-tagged corpus data). We train each baseline model with Adam [KB15], weight decay [LH19], and 0.1 dropout. We also experimented with a slanted triangular learning rate (i.e., using a warm-up and linearly decaying learning

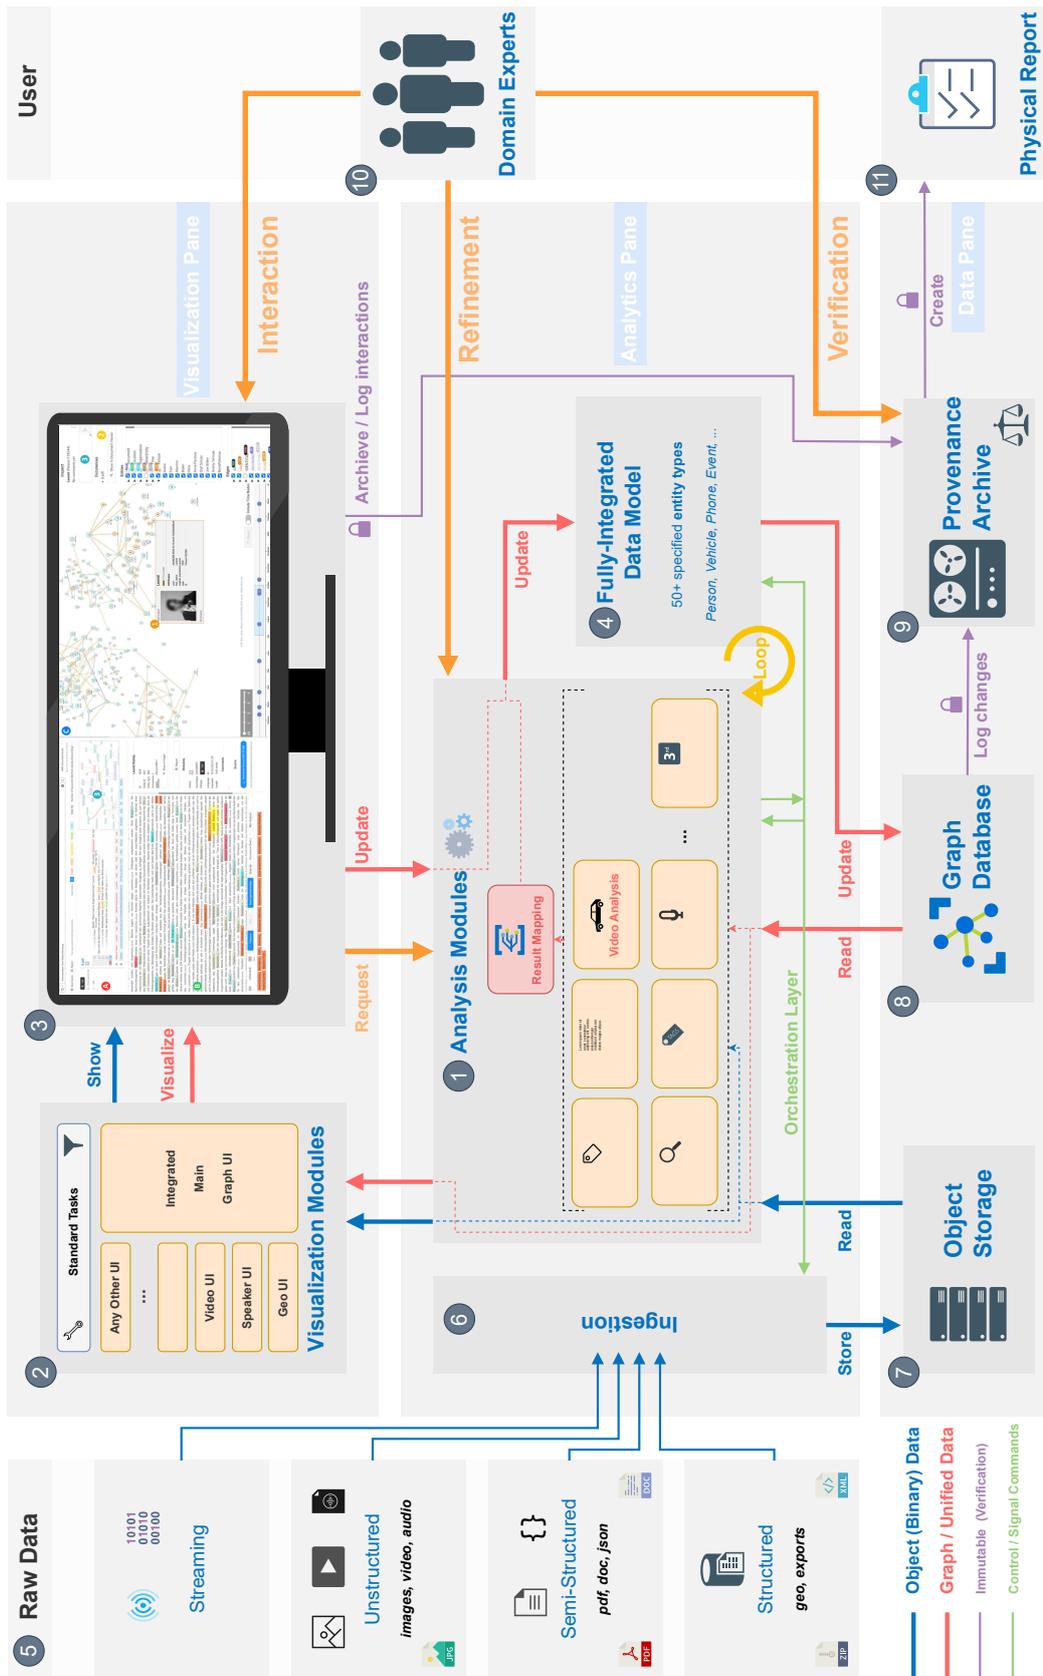


Figure 8.2: High-level architecture of MULTI-CASE, highlighting the main components 1 - 9, like the 1 analysis modules or the 3 graph UI with 2 visualization modules, the 4 fully-integrated graph data model, as well as different data paths and user handling, like the 11 analytics workflow. A detailed description is provided in Section 8.4.

rate) [JS18] and found a slight positive effect on final performance. Early stopping was implemented based on the validation F-score with a number of patience steps of 10. Fine-tuning of all models was performed on a single RTX4000 GPU. We report the full set of hyperparameters and additional results at osf.io/eap4r.

After training, we evaluate the model performance alongside other base models on a held-out test set and describe the results in Section 8.5.4. We also note the recent advances by Large Language Models (LLMs), which can drastically improve specialized NER-tagging through zero- or few-shot learning, in the outlook in Section 8.6.2.

8.4 System Design

The proposed architecture for our framework and the fully-integrated graph data model described in the following is shown in Fig. 8.2. When necessary, we also detail the expert reasoning and the ethical considerations behind individual design decisions while also referring to Sections 8.3 and 8.5.5 as well as 8.5.2 for further discussions on these topics. The guideline numbers for the ethical and privacy reasoning (e.g., C1-6, R1-5, A1-6) refer to the nomenclature established in Chapter 7.

Overall, the system consists of individual plugins ① [Analysis Modules](#) for specific analysis tasks and data types, a ③ [Main Graph-based UI](#) together with specialized ② [Visualization Modules](#) (e.g., text analysis or video-analysis) for a web-based exploration. This fulfills the demand by experts to be capable of specialized analysis that interfaces with an overall case working framework. Similarly, the heavy computations are run on a centralized server, while the interface nowadays is a standard web-based approach running on a regular (or thin) client. One key aspect of the overall system is the ④ [Fully-Integrated Data Model](#) stored in a ⑧ [Graph Database](#), which acts both as a conceptual abstraction layer between modules and a central source of shared knowledge. This enables the experts to work on a consistent data set in an integrated environment and not lose information compared to switching between applications, increasing Efficiency (A4) while addressing the working together of machines and users (C5). Supporting roles fall to the ⑦ [Object Storage](#) to store any input and intermediate data and the ⑨ [Provenance Archive](#) as a [revision-safe storage](#), which is considered essential for Opacity (C3) and Accountability (C6). The ⑩ [domain experts](#) can communicate with the system by interacting with the visualization, forming a collaborative Human-Machine-Configuration (C5), refine the display through analysis parameters, as well as [verify](#) the results, which increases understanding and fosters trust and works against Lack of Accountability (R1), while enabling Human Oversight (R5) and also facilitating a critical reflection (R4). This verification is available both in the interface and in a ⑪ [physical report](#), which the experts still need to document their findings in a structured way.

8.4.1 Data Model

Diverse types of ⑤ **Raw Data** are supported, ranging from unstructured data (images, video, audio), over semi-structured documents (e.g., PDF documents), to structured data types (like geolocation tracks or exports), as well as streaming data. The needs of the domain experts naturally vary here depending on their organization and tasks, but typically the first two types are the most common ones. The input is only limited by the plugin analysis modules. When **data is** ⑥ **ingested**, it is stored in the ⑦ **Object Storage**. Based on the input type, the **Orchestration Layer** selects one or more analysis modules for knowledge extraction, for example, NER for text documents. The main results are mapped to the ④ **Fully-Integrated Data Model** stored in the ⑧ **Graph Database**. For example, for NER, this could be the detected *entities*, like persons, location, or dates, as well as their *relations*, while for video analysis, an object like a car along its properties and a relationship to time and location. Two aspects are of primary importance:

(1) the **data model** ideally has to be as mutually exclusive and collectively exhaustive as possible. The data model was designed with several domain experts and generalized from existing case models like IMP (Information Model Police). In our case, we arrived at 50+ hierarchical *entities* (graph nodes) and 10+ *relationship* types (graph edges), trying to find the right balance between a generic data model and enough specialization. While a very generic data model allows for the reflection of virtually all analysis results, the automatic conclusions, connections, and information enrichment in such a case can remain very limited. In contrast, a highly specialized data model allows to reflect on the findings with high precision and enables many automated conclusions. However, it always poses the danger of being too specified (i.e., available properties on a type) to capture all relevant information. Indeed, the principle design is flexible, subject to change, and can be adapted by adding more specialized entities or data fields. Analysis modules are change-agnostic if the entities and attributes they work with are untouched. In our case, we derived everything from a root *Thing*, with *Entity*, *Event*, *Datetime*, *Location*, and *Document* as the first hierarchical layer, each having further subtypes (e.g., *Person* or *PhoneCall*). For example, a *Timespan*, as a subtype of *Datetime*, represents a specific time range and can be related to a *PhoneCall* via a relationship, which in turn may be related to specific phone numbers, which again might be related as belonging to actual persons. Attributes for each entity store associated information. Through *relationships*, one can also model source attribution (source document and analysis module) and ③ confidence scores, e.g., based on the 6x6 intelligence scoring [UNO11], which many analysts are well familiar with, strengthening Literacy (A5). This scoring can have an influence on automatic decision-making: when certainties are considered by algorithms, this can support working towards Preventing Automated Inequality (R3) and limit Exaggerated Expectations (C4) and

Discriminatory bias (C1) through manual priming. Simultaneously, the opposite could also be true, where the system warns a user of inherent prejudice evident in analysis choices.

(2) The data model allows a structured information **exchange** and also **information enrichment process** between modules, which the experts consider essential. **Updates of the data model** can trigger subsequent runs of other analysis modules when they have signed up for specific creations/updates: for example, an imported audio file might be analyzed first by a speaker detection (with the creation of a specific audio entity), then by a speech-to-text transcription (with a text entity), and then by a NER process, which can result in an **enrichment of the graph** with the conversation content through multiple entities (e.g., persons, location, or times). All changes (creations, updates, hidings) in the graph data are **logged via a write once, read many** ⑨ [Provenance Archive](#).

8.4.2 Component Integration

The individual ① [Analytics Modules](#) like NER or transcription are designed as plugins and can be flexibly combined depending on the analytical needs, allowing for Customization (A6) and ensuring User Agency (A1) of the experts. In this work, we primarily focus on the search and NER modules as an exemplary prototype developed by us, while other modules are provided as open source (e.g., transcription via Whisper [RKX⁺22]) or by commercial partners. During startup, the modules register themselves, their supported data types for ingestion, and the graph change listeners via the [Orchestration Layer](#). Further, each analytics module can register custom *context actions* (e.g., show similar persons) and *preview handlers* (e.g., picture or video player), which are integrated into the ③ [Main Graph UI](#), allowing for a tight coupling between the UI and individual modules functions in ② [specialized UIs](#), supporting the mental mapping of the experts.

8.4.3 Interfaces and Interaction Principles

The interfaces are web-based, and the provided views are **tightly coupled** and **inter-linked**, strengthening the Human-Machine-Configuration (C5) and the User Agency (A1) through Opacity (C3). Entities are consistently mapped via the unified, fully-integrated data model, allowing for the enrichment of information within the main graph-based overview and across views.

The main interface to start explorations is the ③ [Main Graph UI](#) (see also ④ in Fig. 8.1). It provides a highly scalable *GPU-based rendering* of a *Knowledge Graph* (a network-based visualization of the interconnected data items and their relationships), together with several linked views. This graph-based overview is less scalable than a matrix-based approach [FSS⁺21], however, aligns more closely

A *timeline* at the bottom shows both datetime information as part of the Knowledge Graph and document times, allowing for brushing and filtering to optimize the graph view and empower investigators to follow an event- and time-based workflow in alignment with their exploration. When hovering over documents, only these are shown, while selecting zoomable and shiftable ranges restricts the shown parts of the graph. As can be seen from the examples, the amount of information displayed in the graph view is typically quite large, which hampers exploration. Therefore, a ② *Neighborhood Exploration*, acting similar to a magnifying glass or spotlight, allows to show the local neighborhood of a node (for example, 3 or 4 steps), and clicking any visible node transitions to the new neighborhood, allowing for a seamless exploration with a manageable amount of local, contextual information displayed without overloading the users, which can improve Efficiency (A4). Another approach to reducing the amount of clutter is to selectively merge confirmed relations to clusters, for example, aliases for persons or create groups. A slider allows for a confidence level based on the 6x6 system, which means that automated decisions without manual verification are never categorized as (very) likely (A or B), preventing Automated Inequality (R3) and enforcing Fairness (A3) and critical reflection (R4) through Human Oversight (R5).

Due to the amount of information shown (for typical investigations, this can be 30k nodes and 100k edges), we need to use several techniques to achieve 60+ fps performance: The graph is rendered entirely on the GPU and leverages instancing and custom shaders. This results in, once set up, a fixed-sized geometry of a few hundredth vertices and three WebGL draw calls (nodes, edges, labels), resulting in efficient rendering performance. Much of the visualization and visibility status is controlled from within the shaders, with crafted texture atlases and mipmapping for efficient textures, especially for nodes and text labels. To render more than 0.5 million text characters in real-time, we use a pre-generated font texture atlas and supply each node label instance with its correct, fixed-size ASCII-Code label (Unicode would be possible, but increase the texture size). This supply of instance-specific data (e.g., labels, position, node render state) is achieved through uniform buffer objects, acting similarly to a memory map, which is highly efficient. The sidebar uses virtual lists to render on demand, further reducing DOM usage. However, the number of nodes and edges is still limited by JavaScript and Browser performance.

The NER module offers an Ⓐ ontological search and Ⓑ textual view (see Fig. 8.1) as UI components. In the UI, a ① **context overlay** can be shown, for example, over a person's name with a preview image of a person together with other meta-data. This reduces domain boundaries and relieves the mental load of the users. Text understanding can be helped (see Section 8.5.5) by color-coding named entities according to type and offering aggregation and interactions. Linked views at the bottom show all entities in the document grouped by type and ordering, e.g., by

count, can be used to quickly navigate between occurrences through auto-scrolling, highlighting, and stepping.

The **ontological search** uses multiple (de-) selectable semantic search modes (exact match, substring match, fuzzy match, or ontological match). The latter allows searching *semantically* instead of guessing the correct keywords. This ontological search is considered very beneficial by the experts, as it reduces the burden on them to know the exact terms used but more generically describes the concept of what they are looking for. Search results are shown with specific probability scoring based on the distance (steps) taken in an ontology database, linking different properties. One example would be to search for "accommodation" and get results with "hut", "hotel", or "cottage". The quality of the results, of course, depends on the extensiveness of the ontology, which often has to be adapted domain-specifically. Here, the experts can modify the ontology *on the fly*, e.g., to adapt to specific codewords.

Another type of interaction resulting from the tight integration comes even closer to the **traditional visual analytics loop**: While updating analysis parameters within a module usually only affects this module's results, through the fully-integrated data model and module listeners, it becomes possible to achieve *inter*-module exploration and refinement, coming closer to the expected levels of automation by current users. For example, when several speakers in audio files are recognized, and the transcripts are polluted by some of the speakers being background noise, the user can manually deselect the speakers, resynthesizing the audio, and the downstream analysis is automatically re-run, i.e., transcription and then knowledge extraction through NER. Old results can, in this process, be hidden (i.e., flagging the old document and its inference) to avoid an over-cluttering of the graph, which is considered extremely relevant by the experts to allow them to focus on relevant information only but can also be used to preserve Privacy (A2, C2).

While users work with the application, all performed actions are **logged to achieve provenance**, provide Accountability (C6), as well as prevent abuse through Human Oversight (R5).

8.5 Evaluation

We conducted a thorough evaluation of our approach, including feedback from multiple perspectives, to determine the effectiveness of the system. To showcase the practical usefulness of our approach, we present a case study in an investigative journalism setting, supporting war crime investigations (see Section 8.5.1). To scrutinize the ethical and privacy risks involved, we then evaluate our approach based on ethics design guidelines [FHJ⁺22] for intelligence applications (see Section 8.5.2). To judge the resulting capabilities of the developed framework, we use a

state-of-the-art intelligence capability assessment [FDS⁺22b] (see Section 8.5.3). To assess the quality of the underlying language model, we performed benchmarks on relevant NER-task, achieving state-of-the-art performance (see Section 8.5.4). Finally, to evaluate the system from an expert perspective, we conducted a formative user evaluation with eleven domain experts in law enforcement (see Section 8.5.5).

8.5.1 Case Study

In the following, we describe a simplified, *artificial* case study modeled after real-world workflows seen in **investigative journalism**. Here, we describe the process of identifying, placing, attributing, and documenting **war crimes**. We have chosen this example due to its high relevance, the high analysis stakes both for the victims as well as innocent persons, and the plausible availability of large amounts of multimodal data.

Goal — Alisa is an aspiring journalist for the respected newspaper *The Custodian*. She has been reporting about a brutal war in her home country for months now. While there have been some high-profile reports on war atrocities, she knows this is just the tip of the iceberg, and many people are missing. After reading some OSINT (Open Source Intelligence) reports, she wonders if she can also find out more about the forgotten victims of war. Simultaneously, she wants to see the perpetrators held accountable, so she aims to document her findings meticulously and hand her chain of evidence over to the ICC (International Criminal Court), which has started pre-trial investigations.

Data Collection — She starts off by collecting raw data: From various online sources reporting about the war, like Telegram, she exports messages, images, audio, and videos. From a friend and contact working for a large telecommunication provider, she gets a large dump of telephone calls and texts originating from foreign cell phone numbers logged into the telco's network. They were recorded by order of the nation's domestic intelligence agency. Further, on her newspaper's website, she allows for a SecureDrop submission for images and videos. Overall, she ⑤ receives thousands of hours of audio and video and tens of thousands of texts and images, which she imports into MULTI-CASE. The system ingests this data and runs the analysis pipeline.

Initial Exploration — First, Alisa is overwhelmed by the sheer amount of data in the ④ *Knowledge Graph* view. She looks around and randomly starts listening to some recorded phone calls via the ① preview hover menu. Some are hard to understand due to multiple persons talking intermittently in the background.

Analysis Pipeline — The system offers her ① automatically-generated transcripts through the *Speech to Text* module while the audio is played simultaneously. She notices that the transcripts are not perfect when hearing the recordings, but they

still help her a lot, as she can skim over the content in the **B Document Viewer** much quicker. Wondering if the speakers talked about locations, she searched manually for common city names, finding many results. She realizes she can also use the entity search to display all locations the semantic text analysis has found, a summary of which is shown at the bottom. Through these and reading some context, she realizes the transcripts are intermingled with speech fragments (and locations) from the background speakers.

Multimodal Combinations — She **5** jumps back to the graph view and selects the *Speaker Recognition* module for the selected node. It identified four speakers and offered some best shots to listen to, together with individual transcripts. Hearing them in isolation, she realized that two were radio moderators. She deselects both speakers and lets the downstream analysis task run again. In the **C Knowledge Graph** view, the old entry is **4** transparently archived and replaced by the new audio. Now, the recordings and transcripts are much clearer, but listening or reading through only a few would still take hours.

Semantic Search — She decides to **A** *search* literally for some terms and words she suspects might have been used but get fewer results. Instead, she enables the fuzzy as well as the *ontological search*. Now she receives many more results. In some, the spelling seems off, and in others, she gets synonyms and hyponyms for her query. Reading over some of the matching sentences, she realizes several specific words are used and also learns some new ones the system did not detect.

Retraining On-The-Fly — She adds those words to the **4** *built-in ontology* and re-runs the search. As she reads a conversation about a small village where “a lot of — things happened,” she feels she might be on to something. Semantically searching through the remaining transcript in the **B Document Viewer**, the speakers refrain from mentioning the village or such events again.

Cross-Matches — However, the system has recognized the village’s name as a location descriptor and offers her to view it in the **C Knowledge Graph** view. There, she uses the **2** *Neighborhood Exploration* to see all connected entities up to three steps from this town. She finds out that another document mentions this tiny village in a spatial context to a larger town while the village is again allegedly mentioned in connection with some persons named *A* and *B* repeatedly over an extended time period. Using the **C** *timeline view*, she restricts the view to a specific time range where she knows that the area around this larger city was temporarily invaded before the attackers were forced out into the neighboring woods. The graph becomes less crowded, and the system displays a weak link from person *A* to another name *A'* with a longer name form. The weak link comes from yet another transcript, where the persons named are mentioned closely together.

Manual Investigations — Alisa requests her assistant to read the transcript while she briefs her boss about the preliminary findings. After returning, Alisa sees that

her assistant (working collaboratively on the case with her) has concluded that the persons mentioned in the report are likely similar and ③ has assigned a B score (highly likely) for the link in the 6x6 system [UNO11]. The person *A* has also been mentioned in the caption of a Telegram picture. Having used the *Image Analyzer* module, her assistant has found visual matches for this person in several pictures and also two videos, which he has flagged for her. She watches both videos, and one clearly shows a war crime.

Handling Fakes — She also ① searches for *B*, and she finds a graphic image but also sees *B* in a similar setting, seemingly taken weeks prior. She identifies the environment and obtains a broader view of the situation: the image is fake, likely disinformation. She ⑥ adds a comment and marks it as disproved, becoming archived by default.

Progressive Analytics — During her background research, interviewing one ICC representative, she is offered access to the ICC evidence collection platform, where users worldwide can upload materials of suspected war crimes. She also ⑥ imports this potential evidence enriching the underlying data model. Now, she runs a further person search using *Image and Video Analyzer* and finds the picture of a military photo ID. The person in the picture looks very similar to *A*.

Evidence Collection — Using the ⑪ reporting functionality, she prints out a trace of her analysis steps, including the transcripts with reference to the original audio files, the connection network with locations, all the associated imagery and data as a PDF report, and the associated document dump. She plans to hand it over to her ICC contacts and lawyers for them to further verify the potential claims for a subsequent trial. They plan to perform classical investigative work like forensic audio, facial analysis, and site visitation to collect evidence to back up and corroborate the potential war crime she found using the system, now knowing what to look out for.

8.5.2 Ethics Design Guidelines

In Chapter 3, we have discussed in depth the ethical implications of using VA systems in intelligence and derived the first comprehensive overview of *detailed, technical* considerations to take into account when designing such systems. As pointed out, *the ethical implications [have to be considered] as an integral part of the design process from the outset* [FHJ*22]. In the following, we describe how we have applied those considerations during the development of MULTI-CASE. The guideline numbers (e.g., C1-6, R1-5, A1-6) refer to the nomenclature established in Chapter 3.

Semi-automated analyses are used, but the user remains in control for Human Oversight (R5), and the automated decisions are transparent (e.g., through ④ attribution and ③ confidence scoring) for Opacity (C3), addressing User Agency (A1) and

Lack of Accountability (R1). ⑨ Provenance of the analysis steps taken can further strengthen this Human Oversight (R5) and provide Accountability (C6). The ability, for example, to ⑥ flag wrong or unrelated content can support Privacy (A2, C2) aspects by being less intrusive than human verification (as humans might memorize sensitive information). All automated system risk exhibiting inherent Discriminatory Bias (C1), but human operators also do. We published our underlying model for transparency reasons (cf. Opacity (C3)) and to detect or Prevent Automated Inequality (R3). The design as a hybrid Human-Machine-Configurations (C5) inherently ⑤ allows for mutual checks and balances to facilitate more Fairness (A3) and Human Oversight (R5). The semi-automated analysis undoubtedly can ② improve Efficiency (A4), while care was taken not to abstract too much and for the information to remain ① transparently attributable (cf. Opacity (C3), Accountability (C6), Lack of Accountability (R1)), which is achieved through the unified ④ fully-integrated data model. By making clear what aspects are automated and which are manual, by providing ③ confidence scores, and by not offering unrealistic features such as "solve investigation" buttons, one works against Exaggerated Expectation (C4). Effective usage of the system and Literacy (A5) can only come with experience and daily usage. Integrated ⑥ sharing between colleagues, e.g., of saved search filters or information through comments, can support this. However, we note that more could be done here for our approach, but we expect that literacy will primarily be achieved through classical Training and Community-Building Among Users (R2). By enabling ④ interactive modifications to the underlying models like the ontologies, Customization (A6) can help users to adapt the system to their needs. One aspect to further improve upon is automated guidance to facilitate critical reflection (R4), for example, by automatically trying to detect biased behavior by human operators.

8.5.3 Intelligence Capability Assessment

We assess our framework according to a system capabilities classification [FDS⁺22b]. This generic classification aims at knowledge exploration systems, including holistic approaches, focusing on intelligence applications. The classification's main focus is to assess the (technical) capabilities in a structured form, for example, if time-series data is supported, what type of interactions are used, or which type of knowledge is generated through AI support. In this regard, it indirectly includes results from older requirements studies in intelligence [DGLR09; KGS09; KS11; LKT⁺14]. However, these previous studies primarily describe the user interactions with the system like Jigsaw [SGLS07] through Overview and Detail, or Find the Clue and Follow the Trail [KGS09]; those aspects included in the older studies but not in the capability assessment were evaluated as part of the expert evaluation (see Section 8.5.5). We describe and assess our approach according to the 52 criteria

posed in the **classification scheme** [FDS*22b]. The icons indicate ○ no, ◐ partial, and ● full support. For a detailed discussion on the attributes themselves, we refer to the original paper while we provide examples and placement of MULTI-CASE's capabilities in the following:

In the dimension *Data and Information*, MULTI-CASE can compete with the state-of-the-art: It supports all basic **Data types**: *text* like documents and messages, *audio* like recordings, *image/video* like pictures or video recordings, *network* like relationship networks or call records, and ◐ *time-series*, primarily through meta-data like discrete timestamps. Classical, continuous time series are not explicitly supported. Regarding the **coding** of data, only *digital* modalities (i.e., the face-value of information) are supported, not ○ *analogical* ones (e.g., interpretation of facial expressions to detect lies or irony). This is comparable to the vast majority of approaches. Regarding the orthogonal **Expression**, *explicit* information is supported, but also *implicit* one, through the use of the underlying ontologies, which is a rare capability. With regards to communication between **Parties**,  group communications are supported, but not specifically nested groups (i.e., subgroups). Analysis of ○ **Power Relations** is not supported. However, the investigative application is designed in such a way that it accounts for acts of deception and partially considers the ◐ **Measurement Problem**: For example, the use of code words is, in principle, supported through the domain-specific ontologies and specially trained NER model and also by looking at meta-data, which is harder to craft. This is a crucial capability in investigative systems, which many current approaches still delegate fully to the users.

In the dimension *Processing and Models*, our approach is suitable for a wide variety of analyses. Regarding the **Methodologies**, supported are *Representational* analysis to present the information, and especially *Confirmatory* analysis to validate hypothesis as well as *Exploratory* analysis to find relevant, a priori unknown facts. ◐ *Predictive Analytics* is partly integrated, depending on the employed modules. In terms of the employed analytical **Modalities**, all primary ones are equally supported: *Content* like actual text or videos, for example, through the Document Viewer or the Video Analyzer, *Network* for relationship analysis through the Knowledge Graph and Neighborhood Exploration, or *Meta-Data* through the Knowledge Graph and the Timeline in combination with the filtering functions. This holistic, integrated, and interconnected analysis is a crucial factor distinguishing MULTI-CASE from most existing approaches. The **Analysis** itself supports an incremental, streamed data import, making it an *online* analysis. Regarding the **Latency**, the standard use case for an investigative system will be a *delayed*  analysis. One key advantage of the underlying model and the modular architecture is the achieved **Scalability**. It supports huge (III) investigative volumes for *ingress*. Also, through its Neighborhood Exploration, the number of concurrent entries under consideration in the *analysis*

can be regarded as medium (II), more than many other approaches. As our approach is a research prototype and not a commercial application, the support for ○ **Data-Mappings**, like many importers, is limited.

In the dimension *Visual Interface*, many combined strategies are leveraged. Regarding the visualization **Pane**, the usual 2D is supported, but the Knowledge Graph also leverages 3D. Stereoscopic 3D ○ S3D is unsupported but easily addable. Regarding the **Operation Methods** [YKSJ07], all are supported: one can *Select* an entity to get more detailed information from all modules combined, *Explore* different semantic matches or the Knowledge Graph, *Reconfigure* the confidence thresholds for automated merging, *Encode* the data as inferred graph relation representation, *Abstract/Elaborate* by adapting the Neighborhood level or inspect information within a specialized module, *Filter* through the semantic search or a timeline range, and *Connect* by showing graph neighborhoods. The **Manipulation** happens both *directly*, e.g., by selecting entries, and *indirectly*, for example, by choosing specific analysis modes, for example, only showing corroborated cross-matches. The **Goal** of the actions is primarily *data tuning* to show relevant information. However, the approach also partly supports ● *model tuning*, where the interactions influence an underlying mode, e.g., by manually confirming relationships between entities or updating the ontology. The **Strategy** involved in interactions are both *iterative* and *progressive*, which go hand in hand in investigative scenarios. The ● *active learning* depends on the individual analysis modules, through feedback or showing an example.

In the dimension *Knowledge Generation*, the **Explanation** of information is performed through *numerical*, *textual*, and *graphical* representations, for example through scoring/sorting, annotated highlighting, or charting, respectively. The **Transfer Function** operates both on the *machine model*, which updates the underlying model through interactions, as well as the *mental model* of the analysts. **Factors** that are considered in our approach are *confidence*, *trust* and *privacy*. For a more detailed discussion, see Section 8.5.2. With regards to the **Time Dimensionality** of the Knowledge Generation, the approach primarily enables the exploration of past information but also allows conclusions based on this information for the given dataset 🗄️. The **Predictive Power** of the system relates to explaining past events and potential upcoming links but also forms predictions about yet-to-ingest data 🗄️. Regarding the **Evaluations** performed, we present a case study 📖 (see Section 8.5.1), this capability assessment ↔ (see Sections 8.5.3-8.5.2), as well as an expert evaluation 👤 (see Section 8.5.5).

8.5.4 Model Evaluation

We evaluated *our NER model* (Huggingface model via osf.io/eap4r) and five *baseline models* based on a hold-out test set of the re-tagged news dataset that we publish for future benchmarks. Based on preliminary experimental results, we decided to train our own NER classifier based on the weights of the pre-trained GottBERT [STT⁺20] language model. We report precision, recall, and F1-score for each entity label (see Tab. 8.1). As an overall observation, we find that our test dataset is challenging for the NER model, as the achieved performance is below scores reported for existing benchmark datasets like GermEval2014 [BBKP14] for all models, including the baselines.

The best-reported score for the GermEval dataset is 86.8% [STT⁺20] with categories *PERSON*, *LOCATION*, *ORGANISATION* and *MISC*. However, on both our generic news dataset and, in particular, the scenario-specific text data, we see significantly lower performance. Still, our model outperforms or matches baseline models on the core categories while achieving satisfactory performance on most additional categories. Still, we observe a drop in performance in broad, newly introduced categories like *EVENT* or *PRODUCT*.

8.5.5 Domain Expert Evaluation

To showcase the effectiveness of our approach in comparison to existing methods, we conducted an expert evaluation with eleven domain experts (LEA 1-2, RS 1-3, SI 1-3, LAW 1, EE 1, POL 1) working in the context of law enforcement.

Expertise LEA 1 is a recently retired former special police forces commander with a 40-year career, now working as a security consultant for law enforcement agencies. LEA 2 is a leading investigator at a federal police force with a 20-year career investigating organized crime. RS 1 is a research scientist and head of the research department with a 30-year career in speech recognition. RS 2 is a junior researcher and developer working for a federal security agency developing analytical solutions for law enforcement in the area of digital forensics. RS 3 is a junior researcher working for the same federal security agency on the topic of phone analysis. SI 1 is a senior principal research engineer with an almost 30-year career overseeing numerous identity solution projects for an international security company. SI 2 is a project manager with a 25-year career working on video analysis and investigative systems at the same company. SI 3 is a principal research scientist with more than ten years of experience in video object tracking also at this company. LAW 1 is a professor and criminologist specializing in security management, hate crimes, and legal aspects with more than 15 years of experience in the field. EE 1 is a sociologist and ethics advisor offering guidance for security research projects. POL 1

Type	sm			md			lg			BERT-German			XML-RoBERTa			Ours		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
PERSON	.69	.72	.70	.76	.77	.77	.78	.78	.78	.93	.89	.91	.91	.87	.89	.94	.88	.91
ORGANIZATION	.55	.47	.51	.56	.52	.54	.59	.55	.57	.81	.65	.72	.75	.64	.69	.78	.78	.78
LOCATION	.53	.57	.54	.59	.61	.60	.61	.61	.61	.90	.63	.74	.84	.62	.71	.88	.90	.89
MISC (Original)	.14	.29	.19	.17	.37	.24	.18	.36	.24	-	-	-	.30	.45	.36	-	-	-
MISC (Own)	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.15	.22	.18
EVENT	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.99	.40	.57
PRODUCT	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.49	.59	.54
DATETIME	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.99	.99	.99
LANGUAGE	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.98	.95	.96
LAW	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.60	.60	.60
QUANTITY	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.97	.96	.97
NUMBERS	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	.98	.98	.98

Table 8.1: Accuracy on the validation dataset for the five baselines (de_core_news_[sm,md,lg], BERT-GER, XML-RoBERTa) and our NER model

is a project supervisor at a national project management agency overseeing civil security research and policy expert.

Methodology The expert evaluation was conducted as a formative evaluation and took a combined 180 minutes, split into a 60-minute presentation and a 120-minute hands-on evaluation. The 60-minute introduction delivered to all experts described the capabilities of the system on a conceptual level while also demonstrating actions in the form of one to three-minute-long screen recordings. During the evaluation, a single station (27-inch FHD screen, mouse, and keyboard) with the prototype was available to the experts, together with two researchers standing by to help with questions and advice. During this time, one of the experts would typically use the system to explore the prototype while being encouraged to think aloud. The other experts could meanwhile observe, comment, and ask questions. After irregular intervals, the experts switched positions, and usage time between experts varied between five to 20 minutes. During the whole session, the experts were asked questions aligned with a semi-structured interview sheet containing a set of 38 prepared questions covering various aspects of our approach. The session's aim was to elicit the domain experts' opinions about the system and gain insights into how they would use the system in their investigative workflows. Further, the experts were asked to comment on the approaches' capabilities, user-interaction concepts, and visualizations while identifying opportunities for improvements. The detailed findings of this evaluation are presented in the following.

Findings Asked about the **benefits** they see in an investigative framework like MULTI-CASE, the criminal investigators state that they hoped to be relieved of the time-consuming, "extremely high manual workload, which currently requires much personnel" (LEA 1) "and time" (LEA 2). Of course, there are existing use case management systems, but "their usage and the casework is performed very much in a manual way [...] with little technical support" (LEA 1), which becomes a "big problem for mass data" (LEA 2), where "automation can be very helpful" (RS 2). In "particular observations produce very large amounts of video data" (LEA 1). For particular problems, some isolated technical solutions exist at some local partners, for example, geo-based analyses (cf. LEA 1), but access depends on the local support and willingness of the partners to help (cf. LEA 1). Further, one of the most important features for them is to import many different types of multimodal documents like "existing records, images, videos" (LEA 2). Here, MULTI-CASE as a ① "large overview system for multimodal data like audio, text or video has the potential to drastically improve investigative work" (RS 3), making it "uncharted territory" (LEA 1). The other experts strongly agree, noting that currently they lack "a complete picture [in a single system]" (LEA 1) and "nothing in this form exists" (LEA 1): neither for phones (cf. RS 3), speech (cf. LEA 2), or text (cf. LEA 1). "Multimodality is the largest benefit, as everything can be seen in context" (RS 2).

Regarding the **risk of automation**, they are aware of potential pitfalls but do not consider them highly problematic: It is likely that “there are errors in the analysis” (RS 1), for example, by different spellings (cf. RS 1). This, however, can also happen when case workers need “to read through thousands of pages or watch weeks of video recordings, where things might be overlooked and error rates increase with time as frustration increases” (cf. LEA 1). “From an automated perspective, it might not be most important to find everything, but to start and find many relevant things” (POL 1). From a “legal perspective, this might be much more critical, as innocent individuals can become part of an investigation” (RS 1). They note that “automated analysis is less of a problem when there is reasonable suspicion for a suspect, but an infringement on fundamental rights is” (LAW 1). In this regard, the modality differs: “images are considered more critical than voice, which in turn is more critical than text” (cf. LAW 1). Possible ways to solve this are “by not focusing on the subject, but on the right infringements [for involved parties]” (cf. LAW 1). This means automated analysis has the potential to be considered less invasive than manual analysis, but “for example through data economy and short-term storage, but this depends on the case” (EE 1).

From an **ethical perspective**, it might be more “justifiable to let the computer search for targets instead of humans” (EE 1) as the human “remembers” (EE 1) offering potential for misuse, while the system forgets after the comparison. Current approaches “do not consider privacy or ethical aspects sufficiently” (LEA 2) and the investigators are independently responsible on their own to follow the rules - however, “there is a wide gap between theory and practice” (LEA 2). “A verified system that works with high accuracy [and without bias] could be fairer than an arbitrary human” (EE 1), as “many humans are very selective and inherently biased” (LEA 1). Regarding the fear of intransparent, autonomous decisions, it was noted that “the systems are always support systems and humans always the final instance” (LEA 1), and before a “prosecution will always be manually verified” (LEA 1). A problem can arise when “humans become too careless and trust the system too much” (EE 1).

One interesting discussion arose regarding the error rate: From the perspective of an analyst “false positives are less of an issue as they can be manually verified, while false negatives are missed” (LEA 1). From the “perspective of innocents, this is directly inverse, but this again depends on the context” (cf. EE 1). “When misses lead to extreme dangers for others, this can be very bad” (cf. EE 1).

The experts consider ⑥ **collaboration** features relevant, where multiple users can work on the same case, as they sometimes have to work with “widely distributed experts” (cf. LEA 1). Also, the “parallel work between colleagues is nice” (EE 1).

Regarding the central **Knowledge Graph**, many experts agree that it can provide a key overview, as “it is extremely important to show all the relations” (POL 1) and the “connections” (LEA 1), which is a “large advantage” (RS 2). “Showing everything to-

gether is very relevant for keeping an overview" (SI 1). For this, the ② *Neighborhood Exploration* is considered "a must, especially when many data items are loaded" (RS 2), as it allows to reduce the visual clutter and only show contextual information. This is an example of a filtering functionality, which is regarded as "essential" (RS 2). Also, the ability to filter the graph and the mergings by ③ confidence is regarded to be beneficial (RS 1). Similarly, the timeline is also considered "very helpful" (RS 2), as "the time and event sequence is very important for the investigation" (LEA 2). In this context, the interactions are regarded as "very smooth and nice looking" (RS 2). However, some experts questioned "if 3D is necessary" (cf. LEA 1) and would favor the 2D graph that is also available. The graph view can act as a "supportive mental map [...] and a large digital notebook" (LEA 2), which "currently is often only in ones head" (LEA 2). For this, the "comment function" is essential and helpful (cf. LEA 2 and EE 1) to make notes, which can be shared between users. Regarding the confirmatory investigative work, however, the "graph view is less important" (LEA 1), where the "individual analysis modules like the document viewer or audio analysis are more relevant [...] supporting the daily work" (LEA 1). For example, in the document viewer, the "automated recognition of entities in the document which are shown at the bottom with their number of occurrences, is especially helpful, as it allows to get a ④ summary understanding of the content of the text already" (cf. LEA 1). Also, the automated transcription of audio "given sufficient quality, is very important and a key advantage" (cf. LEA 1). Especially relevant is the ability to seamlessly switch between view and modalities, for example, ⑤ "to jump from a node in the graph to the text location in the document viewer" (LEA 1) as well as "jumping to search matches" (LEA 1). However, it was noted by several experts that a proficient usage would "require training" (cf. LEA 1, EE 1, RS 1, LAW 1), after its completion, however, would be a "productivity boost" (EE 1).

In terms of **potential future features**, some ideas were mentioned: Among expected quality-of-life improvements like more file type support (cf. RS 2), one area of improvement could be group conversations (cf. RS 2), for example, through colored attributions also inside the document viewer, the creation of cluster-nodes in the graph view to merge related, but currently less interesting entities (cf. RS 2) or show a modification and usage history from co-workers (cf. RS 3). Also, for the comments and exploration history of colleagues, a "misuse button" (cf. EE 1) would potentially be useful to report incorrect use. Also, some more explainability for the automated parts, i.e., why a "speaker was recognized" (cf. EE 1) as such, would be useful and increase trust. Overall, the approach "will be well usable for semi-automated investigative analysis [...] between a knowledgeable user and a supportive system" (LEA 1).

8.6 Discussion and Future Work

As we demonstrated, our approach enhances the capabilities for multimodal intelligence analytics. In the following, we discuss the valuable lessons we learned during development, the implications of the valuable feedback we received about our prototype, architectural design trade-offs, limitations of the approach, and potential future work that remains.

8.6.1 Findings and Lessons Learned

Based on the evaluations in the previous section, we have succeeded in working towards fulfilling the experts' requirements posed from the beginning: MULTI-CASE is an exemplary centralized, multimodal platform framework that allows several analysts to collaboratively work on cases and empowers users through the transparent inclusion of AI-aided decision-making while relieving them of burdensome tasks and considering ethical design guidelines. Following the UNODC [UNO11] task definitions, the main tasks can be performed: link analysis between entities is supported while also allowing to consider them in the context of the surrounding events based on a timeline. While it supports a basic flow analysis in principle, the visualization modules presented here are not particularly suited for this analysis, but through the modular design, a component operating on the shared data model could be developed. We have seen how the multimodal approach can support the analysis of otherwise difficult-to-detect cross-matches, while a visual analytics-based approach has benefits in terms of agency, accountability, and trust. The experts are open to AI-based solutions, especially when it relieves them of mundane tasks, and they feel supported. Leveraging both computational power and human intuition in a tight feedback loop can positively influence the capabilities of the resulting human-machine configuration. Regarding the displayed results, they tend to believe them at face value to some degree when they seem plausible, somewhat similar to findings reported to them by colleagues.

However, we also saw that experts have high expectations regarding the machine results and—especially when not specifically trained for the system—are rather unforgiving with respect to unexpected or contradictory results. Also, they can be easily annoyed in case they feel the system hinders them, holds them back, or torments them through seemingly obvious confirmations. Based on these observations, we can derive several key findings:

F1: A Holistic Approach Supports Finding Cross-Matches

The case study and expert evaluation shows that intelligence investigations require interconnected, multimodal analytics.

Implication: A holistic approach can combine these different analysis modalities within a single context, reducing domain boundaries and enabling effective search for cross-matches. Especially relevant here is a vertical integration between all analysis modules, for example, through a fully-integrated data model.

F2: Unobtrusive Support-Systems are Accepted

As long as a system remains supportive and unobtrusive, relieving analysts of mundane tasks or providing them with valuable hints and insights on request or through nudging, semi-automated systems are accepted. Tormenting approaches hindering the workflow or being intuitive or unreliable can destroy an initial level of trust placed in the system.

Implication: A self-explanatory, easy-to-use user interface combined with helpful but unobtrusive functions is essential. For this, the right balance has to be found between automation and manual confirmation. Unreliable or inconsistent results (without indications) or hindering of workflows should be strongly avoided

F3: Reduce False Negatives for VA—and False Positives for AI

Initially surprising to us was that for many tasks (e.g., search, filter, linking), the domain experts (both LEA 1-2 and EE 1) prefer the error rate to depend on the automation level: for semi-interactive VA a reduction of false-negatives is often preferable, while automated systems should reduce false-positives.

Implication: Consider the optimization task carefully, as the *cost of error*, where not finding something (i.e., FN) or a wrong attribution (i.e., FP) is considered more costly than the opposite. A missed lead might break the whole investigation, while a wrong attribution might cause serious harm to innocents.

F4: Limited Acceptance of Unreasoned Decisions

At least for now, to support an ethical and privacy-aware analysis and offer transparency, fairness, and accountability while fostering user trust, the experts prefer an explainable, interactive system compared to a fully automated approach.

Implication: Due to the high stakes in this domain, experts have concerns about fully-automated systems that cannot provide a rigorous chain of evidence, which—at least for now—is rarely possible. Future developments might shift this balance.

8.6.2 Limitations and Future Work

Nevertheless, the approach remains a research project with limitations:

The **Knowledge Graph** representation uses custom GPU-optimized rendering achieving excellent performance, but it comes with the disadvantage that some of the more advanced results from graph drawing, like more complex curved lines, are not directly applicable without heavy performance penalties. We also want to highlight that we do not see our contribution in designing state-of-the-art graph drawing but in the interactions, combinations, and linkings between the different modalities for the graph.

The integration of the **underlying language model** itself is modular, such that any other transformer-based NER model can be easily used, as the system features a built-in language detection. However, for the evaluation in this chapter, we only used our customized German NER model due to the domain experts' preferences and expertise. We did not explicitly show a generalization, which we, nevertheless, certainly expect. In the future, off-the-shelf transformer-based NER models can be used, with limitations in the types of detected NER and resulting degradation in relationship inference. Alternative models would need to be fine-tuned with additional NER types, requiring appropriate training data. Another problem in this regard can be the analysis of multi-lingual or inter-lingual text and transcripts.

The recent progress with **Large Language Models (LLMs)** like GPT-4 [Ope23] offers interesting opportunities in this regard. This is, in particular, relevant when models are capable of supporting multiple languages as well as providing up-to-date and case-specific query context, as the *New Bings* underlying Prometheus Model [Rib23] shows to some limited degree. Three domain experts (LEA 1, SI 1, SI 2) in our study tried Chat-GPT on crafted case material and were astonished both by the easy workflow of querying and the (relative) quality of the findings as potential leads. They regard such *text-based, interactive prompting* through LLMs, which imitates basic reasoning and summarization capabilities, as potentially very useful. Integrating such natural language prompts in applications, maybe only in supportive roles, seems very promising. Interestingly, GPT-4 also shows surprising capabilities in zero-shot NER labeling. For testing, we let GPT-4 auto-label a subset of our test data. We achieved this zero-shot labeling by prepending a prompt "Extract named entities of the given types from the following text: person, organization and location", resulting in only slightly less quality than manual, human labeling. This could potentially replace specifically trained NER models, like the one we described in Section 8.3. Recent experiments [GAK23] suggest superior results are possible. While this shows the viability of the transfer learning approach, "close-to-real-life" scenarios often perform worse compared to controlled benchmarks [PUL22]. Therefore, evaluating such scenarios in the wild is important to identify persisting limitations, which can be supported by interactive analysis. Also, care must be taken

to consider the additional risks involved when using LLMs: They do not learn from mistakes outside their limited context window (32k for GPT-4-32k), which is relevant when using all documents as context, and most seriously, they tend to suffer from hallucinations that are hard to detect. Further, employment of such solutions would require on-premise solutions or specialized contracts.

Overall, depending on the jurisdictions, **legal requirements** might regulate the allowed automated analysis tasks [BVe23]. The ethical and privacy-aware design, as well as the semi-automated analysis, always subject to human verification, performed in our approach, should allow for usage even in tightly regulated jurisdictions. The concrete usage in critical cases, however, should be accompanied by a prior legal counsel.

One general limitation in this line of research is the **opaqueness of the intelligence community**. Many systems are classified [Dha17] and capabilities are not shared openly—which can be frustrating from a scientific perspective, hampering progress and introducing problems from ethical and privacy perspectives due to missing accountability. It also remains difficult to recruit domain experts to evaluate and analyze the techniques developed in the scientific community. One way to reduce increased reliance on expert evaluations is to also incorporate general interaction strategy design guidelines derived from numerous user interaction evaluations regarding relevance feedback [KJZ⁺21]. Efforts are ongoing to finance more research in open and accountable intelligence solutions (e.g., within Horizon Europe and others). However, we are well aware that some aspects of this domain will likely remain hidden. With our work, we try to contribute to ongoing research in this domain and discuss ways to make these more accountable.

8.7 Conclusion

Over the last few years, AI-driven models have become increasingly prevalent in many domains. This tendency can also be observed in operational analytics solutions in investigative journalism, intelligence, or law enforcement. These domains, in particular, pose distinct challenges due to their sensitive nature. Two aspects, in particular, stand out: ethical and privacy concerns, as well as difficulties in efficiently combining heterogeneous data sources for multimodal analytics. A lack of such *holistic and multimodal* approaches can lead to biased results and increased manual efforts through domain discontinuities.

To address these two challenges, we present MULTI-CASE, a holistic visual analytics framework that enables the exploration and assessment of heterogeneous information spaces (i.e., unstructured, diverse, and multimodal) supported by an equal joint agency between humans and AI to ensure ethics and privacy awareness. To fulfill these requirements, the system operates on a fully-integrated data model

while featuring type-specific analyses with multiple linked components, including a modality-wide search (i.e., full-text, semantics, and all multimodal analysis results), text, and graph-based analysis. Different information streams are linked in a knowledge graph, providing in-situ explanations and transparent source attributions while facilitating responsible exploration through numerous interlinked explorative modules. We discuss the potential for improvements, for example, in rendering, completeness, or the use of more advanced LLMs.

We demonstrate how our framework fulfills the design goals through state-of-the-art intelligence capability assessments and evaluations according to ethics design guidelines. The underlying transformer model showed state-of-the-art performance on relevant benchmarks. To showcase our prototype's analytical capabilities in practice, we presented a case study describing war crime investigations in the context of investigative journalism. Finally, a formative expert evaluation with eleven domain experts in law enforcement confirms that MULTI-CASE facilitates human agency and steering in security-sensitive, AI-supported analysis processes, addresses ethical and privacy concerns, and provides much-needed analytical capabilities.

With this contribution, we aim to provide more insights into the often opaque workings of the intelligence community and strive towards a more accountable and responsible use of modern AI capabilities.

Part IV | Closing

9

Conclusion

This dissertation discussed communication analysis research from a digital and holistic perspective. This concluding chapter will reflect on our contributions to the field and its implications. We take a retrospective view, summarize our key contributions and place them in a broader context, point to promising research directions in the future, and discuss perspectives on communication analysis.

Contents

9.1 Summary	207
9.2 Future Research Aspects	209
9.3 Concluding Remarks	211

9.1 Summary

The ultimate goal across various economic disciplines when analyzing complex and heterogeneous datasets is to enhance **understanding** of the information contained in the data and the generation of **knowledge** from it. Leveraging computer-aided data analysis techniques in this process primarily addresses the issues of computational power (the amount of data that can be analyzed in a given time) and accuracy (in the sense of reproducible, tireless correctness). At the same time, humans so far are often still superior in areas such as creativity or decision-making with limited information through intuition. Visualization plays a crucial role in this endeavor, as it is essential to communicate the results to the human operators and allow them to show and explain intermediate results and context to them. Interactive visualizations—particularly visual analytics, where the interactive visualizations are integrated as part of the analysis loop—aim to bridge the gap between computers and humans to form an effective joint agency. Ideally, both parties form a symbiotic relationship, mutually empowering each other, so they can achieve more than on

their own. This is especially relevant for tasks that are not clear-cut but more open-ended with an a priori unclear result type, like in communication analysis.

Human communication is often subjective, constrained by social and historical norms, heavily context-dependent, and employs various channels. It is one of the most fundamental ways the human mind expresses creativity and identity, reflecting the multicultural and diverse human society. While many written or spoken languages have rules like spelling or grammar, communication semantics and the information exchange underlying communication are highly complex. Humans can often navigate this complex space effortlessly, while computers and algorithms struggle with these built-in assumptions and uncertainty. Leveraging the power of visual analytics for big data applications in communications is therefore not only beneficial but imperative when one wants to avoid a narrow and isolated analysis. Due to the heterogeneity and difficulty of communication analysis, the problem is addressed in several fields and methods. Natural Language Processing (NLP) works on the textual (content) level, while social network analysis (SNA) looks at the social context and the information flow. Further, metadata analysis aims to learn from patterns in associated metadata, while time-series analysis investigates temporal correlations. However, as described in [Chapter 1](#) and [Chapter 2](#), the existing research is relatively segregated and isolated, missing out on opportunities from a combined analysis. Further, while communication behavior is studied in many different disciplines, much of this research is not considered part of algorithmic solutions. As such, the interdisciplinary analysis has yet to be compiled and considered in a coherent framework in the context of digital communication and its semi-automated analysis in computer science, which is the focus of this dissertation.

The **contributions** of this dissertation consist of four parts: After the introduction, the Part I sets the stage. In [Chapter 2](#), we have positioned our work within the context of related work and compared it with the state-of-the-art. Based on this analysis, we have contributed a formalization of the field through a conceptual framework and outlined the research gaps, most of which have been addressed in the following chapters. Before we then start with the technical aspects of how communication analysis can be conducted in a broader and ethically-ware context, in [Chapter 3](#) we have investigated ethical and privacy considerations of communication analysis and proposed ways in which visual analytics can address them, supporting communication analysis. In Part II, we have discussed different techniques that relate to the identification and interpretation of communication as a first step. In [Chapter 4](#), we first have surveyed hypergraph model visualizations before we presented HYPER-MATRIX, a technique for the identification of communication in [Chapter 5](#). After we have identified possible communication topics and participants, we have focused on their interpretation through metadata pattern analysis

in Chapter 6, describing the concept of Conversational Dynamics. In Part III, we have combined the lessons from previous chapters to discuss holistic approaches. We first focused on a primarily text-based level, where we have presented COMMAID in Chapter 7, which usages a multi-level matrix-based view to analyze communication. We then have taken a more multimodal perspective in Chapter 8, where we presented MULTI-CASE as a knowledge-graph-based multimodal framework for interdisciplinary communication analysis. In this final Part IV and Chapter 9, we have summarized our findings in a broader context, pointed out promising future research aspects, and concluded the dissertation.

9.2 Future Research Aspects

At the end of each previous chapter, this work has already outlined challenges and potential future work. In most cases, these considerations address the techniques and methods we discussed in the individual chapters. We can, however, notice some common topics and overarching challenges spanning the whole research field addressed within this dissertation. In the following paragraphs, we highlight *general* limitations and promising future research challenges and perspectives of visual communication analysis in a broader context.

Analogical Code and Broader Context — The analogical code as defined by Watzlawick et al. [WBJ74] is concerned with the *soft* factors of communication, which refer to the implicit meanings. When one begins to include non-text-based information like audio or video, there can be non-verbal signals such as intonation, body language, facial expressions, or gestures, compared to the written word. These codes can transport and contain important cues that might support the analysis of the meaning and provide additional information about the context. In principle, this type of analysis can be conducted as a module in MULTI-CASE (see Chapter 8) and then provide the results towards the combined analysis framework. Similarly, it can be beneficial to also leverage additional context provided in the analysis, for example, information about the power relations of participants, which can influence aspects like formality, choice of words, irony, or content. This might enable the resolution of contradictory information or ambiguities and lead to a richer analysis.

Scalability — The techniques we proposed and developed as part of this dissertation work well for medium-sized datasets in the order of (tens) of thousands of communication parties and events. The primary limitation is not the amount of information in itself, which can often be managed through adequate steps in the analysis process like pre-filtering, indexing, or aggregation. With sufficient computing resources, our approaches could, in theory, handle many millions of such events without compli-

cations. Many of them, like MULTI-CASE follow a modular and flexible streaming architecture, which load and queries data on demand and, given a powerful data structure behind it, is quite scalable. Also, when filtering, for example, by time ranges and topics is one of the first analysis steps as in COMMAID, the scalability is equally high. The main problem in scalability arises when the potentially interesting amount of possibly correlated information cannot be filtered or reduced through (semi-) automatic means. This is a fundamental problem, as finding cross-matches, in general, has an order of $\mathcal{O}(n^2)$ time complexity when the comparison function is arbitrary. One way around this issue would be to use a well-defined, limited set of comparison metrics, thereby enabling – through pre-computations, sorting, and hash-maps—an $\mathcal{O}(n \log n)$ time complexity. This, however, limits the possible comparison techniques. Secondly, displaying the filtered and cross-matched data in a visual interface is further limited by the available visual space. Even matrix-based visualizations are limited by the theoretical pixel-dense maximum. Another option would then be again to use aggregation, summarization, or intelligent zooming/scrolling, all of which can shift the boundaries further but also have drawbacks as they limit overview and place a burden on the analysts to explore the whole space sufficiently. Opportunities range from the development of pre-computable, relatively generic comparison metrics to further study which information aggregation levels—beyond those presented in HYPER-MATRIX, COMMAID, and MULTI-CASE—are tolerated and accepted by the domain experts.

Interactive Querying with Large Language Models (LLMs) – The recent advancements with extremely large deep learning models based on a transformer architecture, more precisely *Large Language Models (LLMs)*, particularly with models such as GPT-4 [Ope23], have unveiled a plethora of intriguing possibilities [KW23] to explore. Due to their design as a Mixture of Experts (MoE), they demonstrate strong capabilities in multiple domains, especially when working with enough context. Communication analysis can provide these models with enormous amounts, making applying LLMs to communication analysis particularly promising. Current models are still somewhat limited by their finite context window, but with increasing window size or the ability to externalize the context window, these issues become less and less. On the one hand, in the case of multilingual models, translation tasks become less relevant. On the other hand, and more importantly, the inclusion of language models can fundamentally change the way analysts interact with the system: Instead of visually exploring prepared analysis results and tweaking the analysis, the users can use interactive querying of the content. For example, one could ask about providing a summary of a long conversation or about any abnormalities as defined by their usual topics. As we showed in [Chapter 8](#), such an interactive prompting imitates rudimentary reasoning and summarization capabilities and is regarded

as a potential game-changer by domain experts. It is important to note that these prompts are unlikely to answer very complex questions or replace the analysts completely. Instead, they help analysts articulate their intent more naturally instead of needing to mentally translate their intent into algorithmic settings. As such, these techniques can take a supportive role to perform more loosely defined analysis tasks, which can still increase overall performance, explainability, and foster user trust. This becomes especially relevant when the model result generation is a multi-step process, where the model refines and explains itself, similar to repeated prompt invocations with queries to check and explain and refine the previous answer until it converges. Despite these advances, it is essential to acknowledge and study that the transfer learning approach can fail in more complex, real-life scenarios, and LLMs are prone to hallucinations or otherwise erroneous outputs. Therefore, it is extremely relevant to identify and evaluate persisting limitations, while the combination of visual analytics and LLMs can also act as a mutually supervising and supporting joint agency.

9.3 Concluding Remarks

In our quest for understanding communication, this dissertation has addressed the question of how *human communication analysis* can be performed by *digital means* through visual analytics and how the expertise of subject matter experts to gain useful knowledge can be included. By doing so, we faced the issue that in contrast to classical communication analysis, as defined in domains such as social psychology, many existing approaches follow insular and segregated analysis strategies. We, therefore, contributed a *conceptual framework* for communication analysis in the digital age and identified visual analytics as a means to tackle the holistic integration of both different analysis modalities as well as leveraging human expertise (RQ1, see page 22). By doing so, given the delicate topic and its potential misuse, we put a particular emphasis on the ethical and privacy dilemmas arising in large-scale digital communication analysis to shed light on the difficulties and dilemmas faced. From there, we work out the benefits of *ethical-aware approaches* and how visual analytics can contribute to such an aim by taking the interface as a starting point, not only in a technical but rather in a socio-technical way (RQ2). We then continue to focus on *technical solutions and approaches*, discussing in depth the details of identifying and interpreting communication. The lessons learned are applied and described as part of two holistic analysis frameworks, which are evaluated in detail (RQ3). Potential *applications* of this overall work can be found in many intelligence areas, for example, investigative journalism, criminal investigations, or business intelligence, while many of the individual contributions presented as part of this dissertation can be applied more broadly.

Bibliography

- [AA06] Natalia Andrienko and Gennady Andrienko. *Exploratory Analysis of Spatial and Temporal Data*. Berlin/Heidelberg: Springer-Verlag, 2006. ISBN: 3-540-25994-5. DOI: [10.1007/3-540-31190-4](https://doi.org/10.1007/3-540-31190-4).
- [AAG19] Florian M. Artinger, Sabrina Artinger, and Gerd Gigerenzer. “C. Y. A.: Frequency and Causes of Defensive Decisions in Public Administration”. In: *Business Research* 12.1 (2019), pp. 9–25. ISSN: 2198-2627. DOI: [10.1007/s40685-018-0074-2](https://doi.org/10.1007/s40685-018-0074-2).
- [AB20] Shivam Agarwal and Fabian Beck. “Set Streams: Visual Exploration of Dynamic Overlapping Sets”. In: *Computer Graphics Forum* 39.3 (2020), pp. 383–391. ISSN: 01677055. DOI: [10.1111/cgf.13988](https://doi.org/10.1111/cgf.13988).
- [AC18] Mike Ananny and Kate Crawford. “Seeing Without Knowing: Limitations of the Transparency Ideal and its Application to Algorithmic Accountability”. In: *New Media & Society* 20.3 (2018), pp. 973–989. ISSN: 1461-4448. DOI: [10.1177/1461444816676645](https://doi.org/10.1177/1461444816676645).
- [ACK*20] Alessia Antelmi, Gennaro Cordasco, Bogumil Kaminski, Paweł Pralat, Vittorio Scarano, Carmine Spagnuolo, and Przemysław Szufel. “Analyzing, Exploring, and Visualizing Complex Networks via Hypergraphs using SimpleHypergraphs.jl”. In: *Internet Mathematics* 1.1 (2020). DOI: [10.24166/im.01.2020](https://doi.org/10.24166/im.01.2020).
- [ADP*22] Kiana Alikhademi, Emma Drobina, Diandra Prioleau, Brianna Richardson, Duncan Purves, and Juan E. Gilbert. “A Review of Predictive Policing from the Perspective of Fairness”. In: *Artificial Intelligence and Law* 30.1 (2022), pp. 1–17. ISSN: 0924-8463. DOI: [10.1007/s10506-021-09286-4](https://doi.org/10.1007/s10506-021-09286-4).
- [AGV19] Ciano Aydin, Margoth González Woge, and Peter-Paul Verbeek. “Technological Environmentality: Conceptualizing Technology as a Mediating Milieu”. In: *Philosophy & Technology* 32.2 (2019), pp. 321–338. ISSN: 2210-5433. DOI: [10.1007/s13347-018-0309-3](https://doi.org/10.1007/s13347-018-0309-3).
- [AI 20] AI Ethics Impact Group. *From Principles to Practice An Interdisciplinary Framework to Operationalise AI Ethics*. 2020. URL: https://www.berstelmann-stiftung.de/fileadmin/files/BSt/Publikationen/GrauePublikationen/WKIO_2020_final.pdf.
- [AMA*14] Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter J. Rodgers. “Visualizing Sets and Set-typed Data: State-of-the-Art and Future Challenges”. In: *Proceedings of the Eurographics Conference on Visualization (EuroVis)*. 2014. DOI: [10.2312/eurovisstar.20141170](https://doi.org/10.2312/eurovisstar.20141170).

- [AMA*16] Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter J. Rodgers. “The State-of-the-Art of Set Visualization”. In: *Computer Graphics Forum* 35.1 (2016), pp. 234–260. ISSN: 01677055. DOI: **10.1111/cgf.12722**.
- [Ame15] Morgan G. Ames. “Charismatic Technology”. In: *Aarhus Series on Human Centered Computing* 1.1 (2015), p. 12. ISSN: 2445-7221. DOI: **10.7146/aahcc.v1i1.21199**. (Visited on July 13, 2021).
- [Ame16] American Civil Liberties Union. *Statement of Concern About Predictive Policing by ACLU and 16 Civil Rights Privacy, Racial Justice, and Technology Organizations*. 2016. URL: **<https://www.aclu.org/other/statement-concern-about-predictive-policing-aclu-and-16-civil-rights-privacy-racial-justice>**.
- [Ame19] Morgan G. Ames. *The Charisma Machine: The Life, Death, and Legacy of One Laptop per Child*. Infrastructures. Cambridge, MA, USA: MIT Press, 2019. ISBN: 978-0-262-53744-5.
- [Amo14] Louise Amoore. “Security and the Claim to Privacy”. In: *International political sociology* 8.1 (2014), pp. 108–112. URL: **<http://dro.dur.ac.uk/14920/>**.
- [Amo19] Louise Amoore. “Doubt and the Algorithm: On the Partial Accounts of Machine Learning”. In: *Theory, Culture & Society* 36.6 (2019), pp. 147–169. ISSN: 0263-2764. DOI: **10.1177/0263276419851846**.
- [APS14] Jae-wook Ahn, Catherine Plaisant, and Ben Shneiderman. “A Task Taxonomy for Network Evolution Analysis”. In: *IEEE Transactions on Visualization and Computer Graphics* 20.3 (2014), pp. 365–376. DOI: **10.1109/TVCG.2013.238**.
- [Are60] Hannah Arendt. *Vita activa oder Vom tätigen Leben*. Stuttgart: W. Kohlhammer, 1960.
- [ARW19a] Devanshu Arya, Stevan Rudinac, and Marcel Worring. “HyperLearn: A Distributed Approach for Representation Learning in Datasets With Many Modalities”. In: *Proceedings of the 27th ACM International Conference on Multimedia*. 2019, pp. 2245–2253. DOI: **10.1145/3343031.3350572**.
- [ARW19b] Devanshu Arya, Stevan Rudinac, and Marcel Worring. “Predicting Behavioural Patterns in Discussion Forums Using Deep Learning on Hypergraphs”. In: *International Conference on Content-Based Multimedia Indexing (CBMI)*. 2019, pp. 1–6. DOI: **10.1109/CBMI.2019.8877384**.
- [Asa19] Peter M. Asaro. “AI Ethics in Predictive Policing: From Models of Threat to an Ethics of Care”. In: *IEEE Technology and Society Magazine* 38.2 (2019), pp. 40–53. ISSN: 0278-0097. DOI: **10.1109/MTS.2019.2915154**.
- [ASW12] Daniel Angus, Andrew Smith, and Janet Wiles. “Conceptual Recurrence Plots: Revealing Patterns in Human Discourse”. In: *IEEE Transactions on Visualization and Computer Graphics* 18.6 (2012), pp. 988–997. DOI: **10.1109/TVCG.2011.100**.

- [ATPL16] Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preoțiuc-Pietro, and Vasileios Lampos. “Predicting Judicial Decisions of the European Court of Human Rights: A Natural Language Processing Perspective”. In: *PeerJ Computer Science* 2 (2016). doi: [10.7717/peerj-cs.93](https://doi.org/10.7717/peerj-cs.93).
- [Aus12] Jef Ausloos. “The ‘Right to be Forgotten’ – Worth remembering?” In: *Computer Law & Security Review* 28.2 (2012), pp. 143–152. ISSN: 0267-3649. doi: [10.1016/j.clsr.2012.01.006](https://doi.org/10.1016/j.clsr.2012.01.006).
- [AW10] Charu C. Aggarwal and Haixun Wang. “A Survey of Clustering Algorithms for Graph Data”. In: *Managing and mining graph data*. Ed. by Charu C. Aggarwal and Haixun Wang. Vol. 40. Advances in database systems, 1386-2944. New York: Springer, 2010, pp. 275–301. ISBN: 978-1-4419-6044-3. doi: [10.1007/978-1-4419-6045-0_9](https://doi.org/10.1007/978-1-4419-6045-0_9).
- [AW18] Devanshu Arya and Marcel Worring. “Exploiting Relational Information in Social Networks Using Geometric Deep Learning on Hypergraphs”. In: *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*. 2018, pp. 117–125. doi: [10.1145/3206025.3206062](https://doi.org/10.1145/3206025.3206062).
- [Bar14] Karen Barad. “Diffracting Diffraction: Cutting Together-Apart”. In: *Parallax* 20.3 (2014), pp. 168–187. ISSN: 1353-4645. doi: [10.1080/13534645.2014.927623](https://doi.org/10.1080/13534645.2014.927623).
- [Bau15] Eric P.S. Baumer. “Reflective Informatics: Conceptual Dimensions for Designing Technologies of Reflection”. In: *Proceedings of the 33rd CHI Conference on Human Factors in Computing Systems*. Ed. by Bo Begole, Jinwoo Kim, Kori Inkpen, and Woontack Woo. ACM, 2015, pp. 585–594. ISBN: 978-1-4503-3145-6. doi: [10.1145/2702123.2702234](https://doi.org/10.1145/2702123.2702234).
- [Bav50] Alex Bavelas. “Communication Patterns in Task-oriented Groups”. In: *The Journal of the Acoustical Society of America* 22.6 (1950), pp. 725–730. doi: [10.1121/1.1906679](https://doi.org/10.1121/1.1906679).
- [BBDW17] Fabian Beck, Michael Burch, Stephan Diehl, and Daniel Weiskopf. “A Taxonomy and Survey of Dynamic Graph Visualization”. In: *Computer Graphics Forum* 36.1 (2017), pp. 133–159. ISSN: 01677055. doi: [10.1111/cgf.12791](https://doi.org/10.1111/cgf.12791).
- [BBH⁺16] Michael Behrisch, Benjamin Bach, Nathalie Henry Riche, Tobias Schreck, and Jean-Daniel Fekete. “Matrix Reordering Methods for Table and Network Visualization”. In: *Computer Graphics Forum* 35.3 (2016), pp. 693–716. ISSN: 01677055. doi: [10.1111/cgf.12935](https://doi.org/10.1111/cgf.12935).
- [BBKP14] Darina Benikova, Chris Biemann, Max Kisselew, and Sebastian Pado. “Germeval 2014 Named Entity Recognition Shared Task: Companion Paper”. In: *Workshop Proceedings of the 12th edition of the KONVENS conference*. 2014, pp. 104–112.
- [BBL⁺17] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. “Geometric Deep Learning: Going Beyond Euclidean Data”. In: *IEEE Signal Processing Magazine* 34.4 (2017), pp. 18–42. doi: [10.1109/MSP.2017.2693418](https://doi.org/10.1109/MSP.2017.2693418).

- [BDF⁺14] Michael Behrisch, James Davey, Fabian Fischer, Olivier Thonnard, Tobias Schreck, Daniel A. Keim, and Jörn Kohlhammer. “Visual Analysis of Sets of Heterogeneous Matrices Using Projection-Based Distance Functions and Semantic Zoom”. In: *Computer Graphics Forum*. Vol. 33. 2014, pp. 411–420. DOI: [10.1111/cgf.12397](https://doi.org/10.1111/cgf.12397).
- [BDG⁺19] Meredith Broussard, Nicholas Diakopoulos, Andrea L. Guzman, Rediet Abebe, Michel Dupagne, and Ching-Hua Chuan. “Artificial Intelligence and Journalism”. In: *Journalism & Mass Communication Quarterly* 96.3 (2019), pp. 673–695. ISSN: 1077-6990. DOI: [10.1177/1077699019859901](https://doi.org/10.1177/1077699019859901).
- [BE01] François Bertault and Peter Eades. “Drawing Hypergraphs in the Subset Standard (Short Demo Paper)”. In: *Graph Drawing*. Ed. by Joe Marks. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 164–169. ISBN: 978-3-540-44541-8.
- [Ben19] Ruha Benjamin. *Race After Technology: Abolitionist Tools for the New Jim Code*. Medford MA: Polity, 2019. ISBN: 1509526390.
- [Ber52] Bernard Berelson. *Content Analysis in Communication Research*. Glencoe, Ill. : Free Press, 1952. ISBN: 9780028412108.
- [Ber84] Claude Berge. *Hypergraphs: Combinatorics of Finite Sets*. Vol. 45. Elsevier, 1984. ISBN: 978-0-4448-7489-4.
- [BHD⁺15] Benjamin Bach, N. Henry-Riche, T. Dwyer, T. Madhyastha, Jean-Daniel Fekete, and T. Grabowski. “Small MultiPiles: Piling Time to Explore Temporal Patterns in Dynamic Networks”. In: *Computer Graphics Forum* 34.3 (2015), pp. 31–40. ISSN: 01677055. DOI: [10.1111/cgf.12615](https://doi.org/10.1111/cgf.12615).
- [BHJ09] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. “Gephi: An Open Source Software for Exploring and Manipulating Networks”. In: *Proc. of the Int. AAAI Conference on Weblogs and Social Media (ICWSM)*. Vol. 3. 2009, pp. 361–362. DOI: [10.1609/icwsml.v3i1.13937](https://doi.org/10.1609/icwsml.v3i1.13937).
- [BHN17] Solon Barocas, Moritz Hardt, and Arvind Narayanan. “Fairness and Machine Learning”. In: *Advances in Neural Information Processing (NIPS) Tutorial 1* (2017). URL: <https://fairmlbook.org/tutorial1.html>.
- [BISM14] Matthew Brehmer, Stephen Ingram, Jonathan Stray, and Tamara Munzner. “Overview: The Design, Adoption, and Analysis of a Visual Document Mining Tool for Investigative Journalists”. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2271–2280. DOI: [10.1109/TVCG.2014.2346431](https://doi.org/10.1109/TVCG.2014.2346431).
- [BK21] Jascha Bareis and Christian Katzenbach. “Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics”. In: *Science, Technology & Human Values* 47.5 (2021), pp. 855–881. ISSN: 0162-2439. DOI: [10.1177/016224392111030007](https://doi.org/10.1177/016224392111030007).

- [BM02] Vladimir Batagelj and Andrej Mrvar. "Pajek - Analysis and Visualization of Large Networks". In: *Graph Drawing*. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2002, pp. 77–103. doi: [10.1007/3-540-45848-4_54](https://doi.org/10.1007/3-540-45848-4_54).
- [BM19] Sérgio Barbosa and Stefania Milan. "Do Not Harm in Private Chat Apps: Ethical Issues for Research on and with WhatsApp". In: *Westminster Papers in Communication and Culture* 14.1 (2019), pp. 49–65. doi: [10.16997/wpcc.313](https://doi.org/10.16997/wpcc.313).
- [BM98] Vladimir Batagelj and Andrej Mrvar. "Pajek - Program for Large Network Analysis". In: *Connections* 21.2 (1998), pp. 47–57. URL: <http://mrvar.fdv.uni-lj.si/pajek/>.
- [BMBL09] Stephen P. Borgatti, Ajay Mehra, Daniel J. Brass, and Giuseppe Labianca. "Network Analysis in the Social Sciences". In: *Science* 323.5916 (2009), pp. 892–895. ISSN: 00368075. doi: [10.1126/science.1165821](https://doi.org/10.1126/science.1165821).
- [BN11] Ulrik Brandes and Bobo Nick. "Asymmetric Relations in Longitudinal Social Networks". In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2283–2290. doi: [10.1109/TVCG.2011.169](https://doi.org/10.1109/TVCG.2011.169).
- [BPF14] Benjamin Bach, Emmanuel Pietriga, and Jean-Daniel Fekete. "Graph-Diaries: Animated Transitions and Temporal Navigation for Dynamic Networks". In: *IEEE Transactions on Visualization and Computer Graphics* 20.5 (2014), pp. 740–754. doi: [10.1109/TVCG.2013.254](https://doi.org/10.1109/TVCG.2013.254).
- [BTC⁺10] Jiajun Bu, Shulong Tan, Chun Chen, Can Wang, Hao Wu, Lijun Zhang, and Xiaofei He. "Music Recommendation by Unified Hypergraph: Combining Social Media Information and Music content". In: *Proceedings of the 18th ACM International Conference on Multimedia*. 2010, pp. 391–400. doi: [10.1145/1873951.1874005](https://doi.org/10.1145/1873951.1874005).
- [BVe23] BVerfG. *Urteil des Ersten Senats (16.02.2023), BvR 1547/19, Rn. 1–178, ECLI:DE:BVerfG:2023:rs20230216.1bvr154719*. 2023. URL: [ECLI : DE : BVerfG:2023:rs20230216.1bvr154719](https://ecli.de/BVerfG:2023:rs20230216.1bvr154719).
- [BYW20] Justin Bullock, Matthew M. Young, and Yi-Fan Wang. "Artificial Intelligence, Bureaucratic Form, and Discretion in Public Service". In: *Information Polity* 25.4 (2020), pp. 491–506. ISSN: 15701255. doi: [10.3233/IP-200223](https://doi.org/10.3233/IP-200223).
- [Car99] John M. Carrol. "Five Reasons for Scenario-Based Design". In: *Proceedings of the 32nd Annual Hawaii International Conference on System Sciences*. Ed. by Ralph H. Sprague. IEEE Comput. Soc, 1999, p. 11. ISBN: 0-7695-0001-3. doi: [10.1109/HICSS.1999.772890](https://doi.org/10.1109/HICSS.1999.772890).
- [CGZ⁺19] Yi Chen, Zeli Guan, Rong Zhang, Xiaomin Du, and Yunhai Wang. "A Survey on Visualization Approaches for Exploring Association Relationships in Graph Data". In: *Journal of Visualization* 22.3 (2019), pp. 625–639. ISSN: 1343-8875. doi: [10.1007/s12650-019-00551-y](https://doi.org/10.1007/s12650-019-00551-y).
- [Cic76] Marcus Tullius Cicero. *De oratore: Über den Redner*. Vol. Nr. 6884. Universal-Bibliothek. Stuttgart: Reclam, 1976. ISBN: 3150068843.

- [CKG⁺20] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. “Unsupervised Cross-Lingual Representation Learning at Scale”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (2020), pp. 8440–8451. doi: [10.18653/v1/2020.acl-main.747](https://doi.org/10.18653/v1/2020.acl-main.747).
- [CLS⁺12] Nan Cao, Yu-Ru Lin, Xiaohua Sun, D. Lazer, Shixia Liu, and Huamin Qu. “Whisper: Tracing the Spatiotemporal Process of Information Diffusion in Real Time”. In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2649–2658. doi: [10.1109/TVCG.2012.291](https://doi.org/10.1109/TVCG.2012.291).
- [CLT⁺11] Weiwei Cui, Shixia Liu, Li Tan, Conglei Shi, Yangqiu Song, Zekai J. Gao, Xin Tong, and Huamin Qu. “TextFlow: Towards Better Understanding of Evolving Topics in Text”. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2412–2421. doi: [10.1109/TVCG.2011.239](https://doi.org/10.1109/TVCG.2011.239).
- [CMS99] Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Series in Interactive Technologies. San Francisco, CA: Morgan Kaufmann, 1999. ISBN: 978-1-55860-533-6.
- [Coo09] Charles Horton Cooley. *Social Organization: A Study of the Larger Mind*. New York: Charles Scribner’s Sons, 1909. doi: [10.1037/14788-000](https://doi.org/10.1037/14788-000).
- [Cor19] Michael Correll. “Ethical Dimensions of Visualization Research”. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Ed. by Stephen Brewster, Geraldine Fitzpatrick, Anna Cox, and Vassilis Kostakos. New York, NY, USA: ACM, 2019, pp. 1–13. ISBN: 9781450359702. doi: [10.1145/3290605.3300418](https://doi.org/10.1145/3290605.3300418).
- [Cos20] Sasha Costanza-Chock. *Design Justice: Community-Led Practices to Build the Worlds We Need*. New York: MIT Press, 2020. ISBN: 9780262043458. URL: <https://library.oapen.org/handle/20.500.12657/43542>.
- [CPC09] Christopher Collins, Gerald Penn, and Sheelagh Carpendale. “Bubble Sets: Revealing Set Relations With Isocontours Over Existing Visualizations”. In: *Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1009–1016. ISSN: 1077-2626. doi: [10.1109/TVCG.2009.122](https://doi.org/10.1109/TVCG.2009.122).
- [CSJ⁺18] Matthew Conlen, Sara Stalla, Chelly Jin, Maggie Hendrie, Hillary Mushkin, Santiago Lombeyda, and Scott Davidoff. “Towards Design Principles for Visual Analytics in Operations Contexts”. In: *Proc. of the 36th CHI Conference on Human Factors in Computing Systems*. Ed. by Regan Mandryk, Mark Hancock, Mark Perry, and Anna Cox. ACM, 2018, pp. 1–7. ISBN: 978-1-4503-5620-6. doi: [10.1145/3173574.3173712](https://doi.org/10.1145/3173574.3173712).
- [CSJ⁺20] Eren Cakmak, Udo Schlegel, Dominik Jäckle, Daniel A. Keim, and Tobias Schreck. “Multiscale Snapshots: Visual Analysis of Temporal Summaries in Dynamic Graphs”. In: *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2020), pp. 517–527. doi: [10.1109/TVCG.2020.3030398](https://doi.org/10.1109/TVCG.2020.3030398).

- [CSL⁺16] Nan Cao, Conglei Shi, Sabrina Lin, Jie Lu, Yu-Ru Lin, and Ching-Yung Lin. “TargetVue: Visual Analysis of Anomalous User Behaviors in Online Communication Systems”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 280–289. doi: **10.1109/TVCG.2015.2467196**.
- [Dat20] DataWalk Inc. *DataWalk*. 2020. URL: <https://datawalk.com/>.
- [DFB⁺21] Frederik L. Dennig, Maximilian T. Fischer, Michael Blumenschein, Johannes Fuchs, Daniel A. Keim, and Evanthia Dimara. “ParSetgnostics: Quality Metrics for Parallel Sets”. In: *Computer Graphics Forum* 40.3 (2021), pp. 375–386. ISSN: 01677055. doi: **10.1111/cgf.14314**.
- [DFH⁺21] Erinc Dikici, Maël Fabien, Jan Horinek, Joshua Hughes, Miroslav Janošik, Marek Kovac, Petr Motliceck, Hoang H. Nguyen, Shantipriya Parida, Johan Rohdin, et al. “ROXSD: A Simulated Dataset of Communication in Organized Crime”. In: *ISCA Symposium on Security and Privacy in Speech Communication* (2021), pp. 32–36. doi: **10.21437/SPSC.2021-7**.
- [DGLR09] Jonathan Decker, Alex Godwin, Mark A. Livingston, and Denise Royle. “A Scalable Architecture for Visual Data Exploration”. In: *IEEE Symposium on Visual Analytics Science and Technology*. Ed. by John Stasko and Jarke Jan van Wijk. VAST. Piscataway, NJ, USA: IEEE, 2009, pp. 221–222. ISBN: 978-1-4244-5283-5. doi: **10.1109/VAST.2009.5333451**.
- [Dha17] Mandeep K. Dhami. “A Survey of Intelligence Analysts’ Perceptions of Analytic Tools”. In: *2017 European Intelligence and Security Informatics Conference (EISIC)*. IEEE, 2017, pp. 131–134. ISBN: 978-1-5386-2385-5. doi: **10.1109/EISIC.2017.26**.
- [DvSW12] Kasper Dinkla, Marc J. van Kreveld, Bettina Speckmann, and Michel A. Westenberg. “Kelp Diagrams: Point Set Membership Visualization”. In: *Computer Graphics Forum* 31.3.1 (2012), pp. 875–884. ISSN: 01677055. doi: **10.1111/j.1467-8659.2012.03080.x**.
- [DWM⁺19] Michelle Dowling, Nathan Wycoff, Brian Mayer, John Wenskovitch, Scotland Leman, Leanna House, Nicholas Polys, Chris North, and Peter Hauck. “Interactive Visual Analytics for Sensemaking with Big Text”. In: *Big Data Research* 16 (2019), pp. 49–58. ISSN: 22145796. doi: **10.1016/j.bdr.2019.04.003**.
- [EDG⁺08] Niklas Elmqvist, Thanh-Nghi Do, Howard Goodell, Nathalie Henry, and Jean-Daniel Fekete. “ZAME: Interactive Large-Scale Graph Visualization”. In: *2008 IEEE Pacific Visualization Symposium*. 2008, pp. 215–222. doi: **10.1109/PACIFICVIS.2008.4475479**.
- [EGA⁺16] Mennatallah El-Assady, Valentin Gold, Carmela Acevedo, Christopher Collins, and Daniel A. Keim. “ConToVi: Multi-Party Conversation Exploration using Topic-Space Views”. In: *Computer Graphics Forum* 35 (2016), pp. 431–440. ISSN: 01677055. doi: **10.1111/cgf.12919**.

- [EGJ⁺16] Mennatallah El-Assady, Valentin Gold, Markus John, Thomas Ertl, and Daniel A. Keim. "Visual Text Analytics in Context of Digital Humanities". In: *1st IEEE VIS Workshop on Visualization for the Digital Humanities (Vis4DH) as part of the IEEE VIS 2016*. 2016. URL: <https://scibib.dbvis.de/publications/view/686>.
- [Ell18] Geoffrey Ellis, ed. *Cognitive Biases in Visualizations*. 1st ed. 2018. Cham: Springer International Publishing and Imprint: Springer, 2018. ISBN: 978-3030071042.
- [EPT⁺05] William Elm, Scott Potter, James Tittle, David Woods, Justin Grossman, and Emily Patterson. "Finding Decision Support Requirements for Effective Intelligence Analysis Tools". In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 49.3 (2005), pp. 297–301. ISSN: 2169-5067. DOI: [10.1177/154193120504900318](https://doi.org/10.1177/154193120504900318).
- [ESG⁺17] Mennatallah El-Assady, Rita Sevastjanova, Bela Gipp, Daniel A. Keim, and Christopher Collins. "NEREx: Named-Entity Relationship Exploration in Multi-Party Conversations". In: *Computer Graphics Forum* 36 (2017), pp. 213–225. ISSN: 01677055. DOI: [10.1111/cgf.13181](https://doi.org/10.1111/cgf.13181).
- [ESKC18] Mennatallah El-Assady, Rita Sevastjanova, Daniel A. Keim, and Christopher Collins. "ThreadReconstructor: Modeling Reply-Chains to Untangle Conversational Text through Visual Analytics". In: *Computer Graphics Forum* 37.3 (2018), pp. 351–365. ISSN: 01677055. DOI: [10.1111/cgf.13425](https://doi.org/10.1111/cgf.13425).
- [ESS⁺18] Mennatallah El-Assady, Rita Sevastjanova, Fabian Sperrle, Daniel A. Keim, and Christopher Collins. "Progressive Learning of Topic Modeling Parameters: A Visual Analytics Framework". In: *Transactions on Visualization and Computer Graphics* 24.1 (2018), pp. 382–391. ISSN: 1077-2626. DOI: [10.1109/TVCG.2017.2745080](https://doi.org/10.1109/TVCG.2017.2745080).
- [Eub18] Virginia Eubanks. *Automating Inequality: How High-tech Tools Profile, Police, and Punish the Poor*. Illustrated edition. St Martin's Press, 2018. ISBN: 978-1-250-07431-7.
- [Eur16] European Parliament and the European Council. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation): GDPR*. 2016.
- [Eur21a] European Commission. *Proposal for a Directive of the European Parliament and of the Council on Improving Working Conditions in Platform Work (COM(2021) 762)*. 2021. URL: <https://ec.europa.eu/social/BlobServlet?docId=24992&langId=en>.
- [Eur21b] European Commission. *Proposal for a Regulation of the European Parliament and of the Council Laying down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts (COM/2021/206 final): Artificial Intelligence Act*. 2021. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>.

- [Eur21c] European Commission. *White Paper on Artificial Intelligence - A European Approach to Excellence and Trust*. 2021. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52020DC0065&from=EN>.
- [Eur21d] Europol. *European Union Serious and Organised Crime Threat Assessment (SOCTA): A Corrupting Influence: The Infiltration and Undermining of Europe's Economy and Society by Organised Crime*. Luxembourg, 2021. URL: https://www.europol.europa.eu/cms/sites/default/files/documents/socta2021_1.pdf.
- [FAMK22] Maximilian T. Fischer, Dennis Ackermann, Yannick Metz, and Daniel A. Keim. "Agent-based Visual Analysis of Public Transport Infrastructure". In: *2022 Smart Country Convention (SCCON)*. Berlin, 2022.
- [FAS⁺20] Maximilian T. Fischer, Devanshu Arya, Dirk Streeb, Daniel Seebacher, Daniel A. Keim, and Marcel Worring. "Visual Analytics for Temporal Hypergraph Model Exploration". In: *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2020), pp. 550–560. DOI: [10.1109/TVCG.2020.3030408](https://doi.org/10.1109/TVCG.2020.3030408).
- [FBBP16] Jeffrey Fagan, Anthony A. Braga, Rod K. Brunson, and April Pattavina. "Stops and Stares: Street Stops, Surveillance, and Race in the New Policing". In: *Fordham Urban Law Journal* 43.3 (2016), p. 77. URL: <https://ir.lawnet.fordham.edu/ulj/vol43/iss3/3>.
- [FDS⁺22a] Maximilian Fischer, Frederik Dennig, Daniel Seebacher, Daniel A. Keim, and Mennatallah El-Assady. "Open Data: Communication Analysis Survey". In: (2022). DOI: [10.17605/OSF.IO/UX53S](https://doi.org/10.17605/OSF.IO/UX53S).
- [FDS⁺22b] Maximilian T. Fischer, Frederik Dennig, Daniel Seebacher, Daniel A. Keim, and Mennatallah El-Assady. "Communication Analysis through Visual Analytics: Current Practices, Challenges, and New Frontiers". In: *2022 IEEE Visualization in Data Science (VDS)*. 2022, pp. 6–16. ISBN: 978-1-6654-5721-7. DOI: [10.1109/VDS57266.2022.00006](https://doi.org/10.1109/VDS57266.2022.00006).
- [FFKS21] Maximilian T. Fischer, Alexander Frings, Daniel A. Keim, and Daniel Seebacher. "Towards a Survey on Static and Dynamic Hypergraph Visualizations". In: *2021 IEEE Visualization Conference (VIS)*. IEEE, 2021, pp. 81–85. DOI: [10.1109/VIS49827.2021.9623305](https://doi.org/10.1109/VIS49827.2021.9623305).
- [FG14] Weiguo Fan and Michael D. Gordon. "The power of social media analytics". In: *Communications of the ACM* 57.6 (2014), pp. 74–81. ISSN: 00010782. DOI: [10.1145/2602574](https://doi.org/10.1145/2602574).
- [FHJ⁺22] Maximilian T. Fischer, Simon D. Hirsbrunner, Wolfgang Jentner, Matthias Miller, Daniel A. Keim, and Paula Helm. "Promoting Ethical Awareness in Communication Analysis: Investigating Potentials and Limits of Visual Analytics for Intelligence Applications". In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, 2022, pp. 877–889. ISBN: 978-1-4503-9352-2. DOI: [10.1145/3531146.3533151](https://doi.org/10.1145/3531146.3533151).

- [FHN⁺07] Xiaoyan Fu, Seok-Hee Hong, Nikola S. Nikolov, Xiaobin Shen, Yingxin Wu, and Kai Xu. "Visualization and Analysis of Email Networks". In: *Proceedings of the 6th International Asia-Pacific Symposium on Visualization*. PacificVis. IEEE, 2007, pp. 1–8. DOI: [10.1109/APVIS.2007.329302](https://doi.org/10.1109/APVIS.2007.329302).
- [FIBK17] Johannes Fuchs, Petra Isenberg, Anastasia Bezerianos, and Daniel A. Keim. "A Systematic Review of Experimental Studies on Data Glyphs". In: *IEEE Transactions on Visualization and Computer Graphics* 23.7 (2017), pp. 1863–1879. DOI: [10.1109/TVCG.2016.2549018](https://doi.org/10.1109/TVCG.2016.2549018).
- [Fis18] Maximilian T. Fischer. *Visual Analytics for Detecting Patterns in Large Communication Networks*. 2018.
- [FKS19] Maximilian T. Fischer, Daniel A. Keim, and Manuel Stein. "Video-Based Analysis of Soccer Matches". In: *Proceedings of the 2nd International Workshop on Multimedia Content Analysis in Sports*. MMSports. New York, NY, USA: ACM, 2019, pp. 1–9. DOI: [10.1145/3347318.3355515](https://doi.org/10.1145/3347318.3355515).
- [FL20] Sofie Flensburg and Signe Sophus Lai. "Mapping Digital Communication Systems: Infrastructures, Markets, and Policies as Regulatory Forces". In: *Media, Culture & Society* 42.5 (2020), pp. 692–710. ISSN: 0163-4437. DOI: [10.1177/0163443719876533](https://doi.org/10.1177/0163443719876533).
- [FM08] Peter W. Foltz and Melanie J. Martin. "Automated Communication Analysis of Teams". In: *Team Effectiveness in Complex Organizations*. Routledge, 2008. ISBN: 9780203889312.
- [FMJ⁺24] Maximilian T. Fischer, Yannick Metz, Lucas Joos, Matthias Miller, and Daniel A. Keim. "MULTI-CASE: A Transformer-based Ethics-aware Multimodal Investigative Intelligence Framework". In: (2024), pp. 1–16. DOI: [10.48550/arXiv.2401.01955](https://doi.org/10.48550/arXiv.2401.01955).
- [FNS13] Azadeh Forghani, Carman Neustaedter, and Thecla Schiphorst. "Investigating the Communication Patterns of Distance-Separated Grandparents and Grandchildren". In: *CHI '13 Extended Abstracts on Human Factors in Computing Systems*. Ed. by Wendy E. Mackay, Stephen Brewster, and Susanne Bødker. ACM, 2013, p. 67. ISBN: 9781450319522. DOI: [10.1145/2468356.2468370](https://doi.org/10.1145/2468356.2468370).
- [FPM92] William J. Frawley, Gregory Piatetsky-Shapiro, and Christopher J. Matheus. "Knowledge Discovery in Databases: An Overview". In: *AI Magazine* 13.3 (1992), p. 57. DOI: [10.1609/aimag.v13i3.1011](https://doi.org/10.1609/aimag.v13i3.1011).
- [Fra05] Cortney A. Franklin. "Male Peer Support and the Police Culture". In: *Women & Criminal Justice* 16.3 (2005), pp. 1–25. ISSN: 0897-4454. DOI: [10.1300/J012v16n03_01](https://doi.org/10.1300/J012v16n03_01).
- [FSS⁺21] Maximilian T. Fischer, Daniel Seebacher, Rita Sevastjanova, Daniel A. Keim, and Mennatallah El-Assady. "CommAID: Visual Analytics for Communication Analysis through Interactive Dynamics Modeling". In: *Computer Graphics Forum* 40.3 (2021), pp. 25–36. ISSN: 01677055. DOI: [10.1111/cgf.14286](https://doi.org/10.1111/cgf.14286).

- [FWY⁺18] Siwei Fu, Yong Wang, Yi Yang, Qingqing Bi, Fangzhou Guo, and Huamin Qu. “VisForum: A Visual Analysis System for Exploring User Groups in Online Forums”. In: *ACM Transactions on Interactive Intelligent Systems* 8.1 (2018), pp. 1–21. ISSN: 2160-6455. DOI: **10.1145/3162075**.
- [FYZ⁺19] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. “Hypergraph Neural Networks”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 2019, pp. 3558–3565. DOI: **10.1609/aaai.v33i01.33013558**.
- [FZC⁺18] Siwei Fu, Jian Zhao, Hao Fei Cheng, Haiyi Zhu, and Jennifer Marlow. “T-Cal: Understanding Team Conversation Data with Calendar-based Visualization”. In: *Proc. of the 36th CHI Conference on Human Factors in Computing Systems*. Ed. by Regan Mandryk, Mark Hancock, Mark Perry, and Anna Cox. ACM, 2018, pp. 1–13. ISBN: 978-1-4503-5620-6. DOI: **10.1145/3173574.3174074**.
- [FZCQ17] Siwei Fu, Jian Zhao, Weiwei Cui, and Huamin Qu. “Visual Analysis of MOOC Forums with iForum”. In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2017), pp. 201–210. DOI: **10.1109/TVCG.2016.2598444**.
- [GA08] Laura J. Gurak and Smiljana Antonijevic. “The Psychology of Blogging”. In: *American Behavioral Scientist* 52.1 (2008), pp. 60–68. ISSN: 0002-7642. DOI: **10.1177/0002764208321341**.
- [GAK23] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. *ChatGPT Outperforms Crowd-Workers for Text-Annotation Tasks*. 2023. DOI: **10.48550/arXiv.2303.15056**.
- [Gal12] Alexander R. Galloway. *The Interface Effect*. Wiley, 2012. ISBN: 9780745662534.
- [Gan21] Boaz Ganor. “Artificial or Human: A New Era of Counterterrorism Intelligence?” In: *Studies in Conflict & Terrorism* 44.7 (2021), pp. 605–624. ISSN: 1057-610X. DOI: **10.1080/1057610X.2019.1568815**.
- [Ger69] George Gerbner. *Analysis of Communication Content*. Wiley, 1969. ISBN: 9780471296607.
- [GF16] Clare Garvie and Jonathan Frankle. “Facial-Recognition Software Might Have a Racial Bias Problem”. In: *The Atlantic* (2016). URL: **<https://www.theatlantic.com/technology/archive/2016/04/the-und erlying-bias-of-facial-recognition-systems/476991/>**.
- [GF17] Ge Gao and Susan R. Fussell. “A Kaleidoscope of Languages”. In: *Proceedings of the 35th CHI Conference on Human Factors in Computing Systems*. Ed. by Gloria Mark, Susan Fussell, Cliff Lampe, m.c. schraefel m.c., Juan Pablo Hourcade, Caroline Appert, and Daniel Wigdor. New York: ACM, 2017, pp. 760–772. ISBN: 9781450346559. DOI: **10.1145/3025453.3025839**.

- [GFC05] Mohammad Ghoniem, Jean-Daniel Fekete, and Philippe Castagliola. “On the Readability of Graphs Using Node-link and matrix-based Representations: A Controlled Experiment and Statistical Analysis”. In: *Information Visualization 4.2* (2005), pp. 114–135. ISSN: 1473-8724. DOI: [10.1057/palgrave.ivs.9500092](https://doi.org/10.1057/palgrave.ivs.9500092).
- [GH11] Georg Groh and Jan Hauffa. “Characterizing Social Relations via NLP-based Sentiment Analysis”. In: *Fifth International AAAI Conference on Weblogs and Social Media*. 2011. DOI: [10.1609/icwsm.v5i1.14157](https://doi.org/10.1609/icwsm.v5i1.14157).
- [Git13] Lisa Gitelman. *Raw Data is an Oxymoron*. Infrastructures series. Cambridge, Massachusetts: The MIT Press, 2013. ISBN: 9780262518284.
- [GKL⁺13] Sohaib Ghani, Bum Chul Kwon, Seungyoon Lee, Ji Soo Yi, and Niklas Elmquist. “Visual Analytics for Multimodal Social Network Analysis: A Design Study with Social Scientists”. In: *IEEE Transactions on Visualization and Computer Graphics 19.12* (2013), pp. 2032–2041. DOI: [10.1109/TVCG.2013.223](https://doi.org/10.1109/TVCG.2013.223).
- [GKN17] Robin C. Geyer, Tassilo Klein, and Moin Nabi. *Differentially Private Federated Learning: A Client Level Perspective*. 2017. DOI: [10.48550/arXiv.1712.07557](https://doi.org/10.48550/arXiv.1712.07557).
- [GMT⁺18] Riccardo Guidotti, Anna Monreale, Franco Turini, Dino Pedreschi, and Fosca Giannotti. “A Survey Of Methods For Explaining Black Box Models”. In: *ACM Computing Surveys 51.5* (2018), pp. 1–42. DOI: [10.1145/3236009](https://doi.org/10.1145/3236009).
- [GRE15] Valentin Gold, Christian Rohrdantz, and Mennatallah El-Assady. “Exploratory Text Analysis using Lexical Episode Plots”. In: *Proceedings of the Eurographics Conference on Visualization (EuroVis)*. EuroVis. The Eurographics Association, 2015. DOI: [10.2312/eurovisshort.20151130](https://doi.org/10.2312/eurovisshort.20151130).
- [Gre18] Ben Green. *Putting the J(ustice) in FAT*. 2018. URL: <https://medium.com/berkman-klein-center/putting-the-j-ustice-in-fat-28da2b8eae6d>.
- [GRF08] Tera Marie Green, William Ribarsky, and Brian Fisher. “Visual Analytics for Complex Concepts Using a Human Cognition Model”. In: *2008 IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 2008, pp. 91–98. ISBN: 978-1-4244-2935-6. DOI: [10.1109/VAST.2008.4677361](https://doi.org/10.1109/VAST.2008.4677361).
- [Gro23] Celeste Groenewald. “Reasoning in Criminal intelligence Analysis Through an Argumentation Theory-based Framework”. PhD thesis. Middlesex University, 2023.
- [GS14] David Gotz and Harry Stavropoulos. “DecisionFlow: Visual Analytics for High-Dimensional Temporal Event Sequence Data”. In: *IEEE Transactions on Visualization and Computer Graphics 20.12* (2014), pp. 1783–1792. DOI: [10.1109/TVCG.2014.2346682](https://doi.org/10.1109/TVCG.2014.2346682).

- [GSBO14] Mohammad Ghoniem, Georgiy Shurkhovetsky, Ahmed Bahey, and Benoît Otjacques. “VAFLE: Visual Analytics of Firewall Log Events”. In: *Visualization and Data Analysis 2014*. Ed. by Pak Chung Wong, David L. Kao, Ming C. Hao, and Chaomei Chen. Vol. 9017. SPIE Proceedings. SPIE, 2014. DOI: **10.1117/12.2037790**.
- [GWA⁺17] Celeste Groenewald, B. L. William Wong, Simon Attfeld, Peter Passmore, and Neesha Kodagoda. “How Analysts Think: How Do Criminal Intelligence Analysts Recognise and Manage Significant Information?” In: *2017 European Intelligence and Security Informatics Conference (EISIC)*. IEEE, 2017, pp. 47–53. ISBN: 978-1-5386-2385-5. DOI: **10.1109/EISIC.2017.15**.
- [Ham03] Frank van Ham. “Using Multilevel Call Matrices in Large Software Projects”. In: *IEEE Symposium on Information Visualization 2003*. IEEE, 2003, pp. 227–232. ISBN: 0-7803-8154-8. DOI: **10.1109/INFVIS.2003.1249030**.
- [Har17] Donna J. Haraway. *Staying with the Trouble: Making Kin in the Chthulucene. Experimental Futures*. Illustrated edition. Academic Press, 2017. ISBN: 978-0-8223-6224-1.
- [HB05] Remko Helms and Kees Buijsrogge. “Knowledge Network Analysis: A Technique to Analyze Knowledge Management Bottlenecks in Organizations”. In: *16th International Workshop on Database and Expert Systems Applications (DEXA’05)*. 2005, pp. 410–414. DOI: **10.1109/DEXA.2005.127**.
- [HBS⁺20] Tom Horak, Philip Berger, Heidrun Schumann, Raimund Dachsel, and Christian Tominski. “Responsive Matrix Cells: A Focus+Context Approach for Exploring and Editing Multivariate Graphs”. In: *IEEE Transactions on Visualization and Computer Graphics PP (2020)*. DOI: **10.1109/TVCG.2020.3030371**.
- [HC14a] Benjamin Heintz and Abhishek Chandra. “Beyond Graphs”. In: *ACM SIGMETRICS Performance Evaluation Review 41.4 (2014)*, pp. 94–97. ISSN: 01635999. DOI: **10.1145/2627534.2627563**.
- [HC14b] E. Hoque and G. Carenini. “ConVis: A Visual Text Analytic System for Exploring Blog Conversations”. In: *Computer Graphics Forum 33.3 (2014)*, pp. 221–230. ISSN: 01677055. DOI: **10.1111/cgf.12378**.
- [HC15] Enamul Hoque and Giuseppe Carenini. “ConVisIT: Interactive Topic Modeling for Exploring Asynchronous Online Conversations”. In: *Proceedings of the 20th ACM International Conference on Intelligent User Interfaces*. Ed. by Oliver Brdiczka, Polo Chau, Giuseppe Carenini, Shimei Pan, and Per Ola Kristensson. IUI. New York, USA: ACM, 2015, pp. 169–180. ISBN: 978-1-4503-3306-1. DOI: **10.1145/2678025.2701370**.
- [HC16] Enamul Hoque and Giuseppe Carenini. “MultiConVis: Visual Text Analytics System for Exploring a Collection of Online Conversations”. In: *IUI’16*. Ed. by Jeffrey Nichols, Jalal Mahmud, John O’Donovan, Cristina Conati, and Massimo Zancanaro. New York, USA: ACM, 2016, pp. 96–107. ISBN: 9781450341370. DOI: **10.1145/2856767.2856782**.

- [HDL⁺09] Rachid Hadjidj, Mourad Debbabi, Hakim Lounis, Farkhund Iqbal, Adam Szporer, and Djamel Benredjem. "Towards an Integrated E-mail Forensic Analysis Framework". In: *Digital Investigation* 5.3-4 (2009), pp. 124–137. DOI: [10.1016/j.diin.2009.01.004](https://doi.org/10.1016/j.diin.2009.01.004).
- [Hen91] Kathryn Henderson. "Flexible Sketches and Inflexible Data Bases: Visual Communication, Conscriptio Devices, and Boundary Objects in Design Engineering". In: *Science, Technology & Human Values* 16.4 (1991), pp. 448–473. ISSN: 0162-2439. DOI: [10.1177/016224399101600402](https://doi.org/10.1177/016224399101600402).
- [Her15] Susan C. Herring. *New Frontiers in Interactive Multimodal Communication*. Routledge London, 2015. ISBN: 9780367466459.
- [Her19] Susan C. Herring. "The Coevolution of Computer-Mediated Communication and Computer-Mediated Discourse Analysis". In: *Analyzing Digital Discourse*. Ed. by Patricia Bou-Franch and Pilar Garcés-Conejos Blitvich. Springer, 2019, pp. 25–67. ISBN: 978-3-319-92662-9. DOI: [10.1007/978-3-319-92663-6-2](https://doi.org/10.1007/978-3-319-92663-6-2).
- [HF06] Nathalie Henry and Jean-Daniel Fekete. "MatrixExplorer: A Dual-Representation System to Explore Social Networks". In: *IEEE Transactions on Visualization and Computer Graphics* 12.5 (2006), pp. 677–684. DOI: [10.1109/TVCG.2006.160](https://doi.org/10.1109/TVCG.2006.160).
- [HF07] Nathalie Henry and Jean-Daniel Fekete. "MatLink: Enhanced Matrix Visualization for Analyzing Social Networks". In: *Human-Computer Interaction – INTERACT 2007*. Ed. by Cécilia Baranauskas, Philippe Palanque, Julio Abascal, and Simone Diniz Junqueira Barbosa. Vol. 4663. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 288–302. ISBN: 978-3-540-74799-4. DOI: [10.1007/978-3-540-74800-7_24](https://doi.org/10.1007/978-3-540-74800-7_24).
- [HF10] Nathalie Henry Riche and Jean-Daniel Fekete. "Novel Visualizations and Interactions for Social Networks Exploration". In: *Handbook of Social Network Technologies and Applications*. Ed. by Borko Furht. Springer, 2010. ISBN: 9781441971418.
- [HFM07] Nathalie Henry, Jean-Daniel Fekete, and Michael J. McGuffin. "NodeTrix: A Hybrid Visualization of Social Networks". In: *IEEE Trans. on Visualization and Computer Graphics* 13.6 (2007), pp. 1302–1309. DOI: [10.1109/TVCG.2007.70582](https://doi.org/10.1109/TVCG.2007.70582).
- [HH21] Paula Helm and Thilo Hagendorff. "Beyond the Prediction Paradigm: Challenges for AI in the Struggle Against Organized Crime". In: *Law and Contemporary Problems* 84.3 (2021), pp. 1–17. ISSN: 0023-9186.
- [HM21] Keith J. Hayward and Matthijs M. Maas. "Artificial Intelligence and Crime: A Primer for Criminologists". In: *Crime, Media, Culture: An International Journal* 17.2 (2021), pp. 209–233. ISSN: 1741-6590. DOI: [10.1177/1741659020917434](https://doi.org/10.1177/1741659020917434).
- [Hon19] Matthew Honnibal. *SpaCy: An NLP library*. 2019. URL: <https://www.spacy.io>.

- [Hoo14] Branden Hookway. *Interface*. MIT Press, 2014. ISBN: 9780262322638.
- [HRD⁺15] Y. Han, A. Rozga, N. Dimitrova, G. D. Abowd, and J. Stasko. “Visual Analysis of Proximal Temporal Relationships of Social and Communicative Behaviors”. In: *Computer Graphics Forum* 34.3 (2015), pp. 51–60. ISSN: 01677055. DOI: **10.1111/cgf.12617**.
- [HSCW13] Steffen Hadlak, Heidrun Schumann, Clemens H. Cap, and Till Wollenberg. “Supporting the Visual Analysis of Dynamic Networks by Clustering Associated Temporal Attributes”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2267–2276. DOI: **10.1109/TVCG.2013.198**.
- [HSS⁺] Johannes Häußler, Manuel Stein, Daniel Seebacher, Halldor Janetzko, Tobias Schreck, and Daniel A. Keim. “Visual Analysis of Urban Traffic Data Based on High-Resolution and High-Dimensional Environmental Sensor Data”. In: *Workshop on Visualisation in Environmental Sciences (EnvirVis)*. DOI: **10.2312/envirvis.20181138**.
- [HTKK08] TaeHyun Hwang, Ze Tian, Rui Kuangy, and Jean-Pierre Kocher. “Learning on Weighted Hypergraphs to Integrate Protein Interactions and Gene Expressions for Cancer Outcome Prediction”. In: *2008 Eighth IEEE International Conference on Data Mining*. 2008, pp. 293–302. DOI: **10.1109/ICDM.2008.37**.
- [HvB18] Chris Jay Hoofnagle, Bart van der Sloot, and Frederik Zuiderveen Borgesius. “The European Union General Data Protection Regulation: What it is and What it Means”. In: *Information & Communications Technology Law* 28.1 (2018). DOI: **10.1080/13600834.2019.1573501**.
- [HW13] Gisela Håkansson and Jennie Westander. *Communication in Humans and Other Animals*. Vol. 4. Advances in interaction studies. Amsterdam: John Benjamins Publishing Company, 2013. ISBN: 9789027272010.
- [IBM20] IBM. *i2 Analyst’s Notebook*. 2020. URL: <https://www.ibm.com/us-en/marketplace/analysts-notebook>.
- [JAF10] Deepak Jagdish, Abbas Attarwala, and Ute Fischer. “Behavior assessment and visualization tool”. In: *CHI ’10 Extended Abstracts on Human Factors in Computing Systems*. Ed. by Elizabeth Mynatt, Don Schoner, Geraldine Fitzpatrick, Scott Hudson, Keith Edwards, and Tom Rodden. New York, NY, USA: ACM, 4092010, pp. 3871–3876. ISBN: 9781605589305. DOI: **10.1145/1753846.1754071**.
- [Jas15] Sheila Jasanoff. “Future Imperfect: Science, Technology, and the Imaginations of Modernity”. In: *Dreamscapes of Modernity*. University of Chicago Press, 2015, pp. 1–33. ISBN: 978-0-226-27666-3. DOI: **10.7208/9780226276663-001**.

- [JFDK00] Carlos Jensen, Shelly D. Farnham, Steven M. Drucker, and Peter Kollock. "The Effect of Communication Modality on Cooperation in Online Environments". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '00*. Ed. by Thea Turner and Gerd Szwillus. New York, New York, USA: ACM Press, 2000, pp. 470–477. ISBN: 1581132166. DOI: [10.1145/332040.332478](https://doi.org/10.1145/332040.332478).
- [JFKF23] Lucas Joos, Maximilian T. Fischer, Daniel A. Keim, and Johannes Fuchs. "Aesthetic-Driven Navigation for Node-Link Diagrams in VR". In: *Proceedings of the 2023 ACM Symposium on Spatial User Interaction (SUI)*. Vol. 25. Association for Computing Machinery, 2023. DOI: [10.1145/3607822.3614537](https://doi.org/10.1145/3607822.3614537).
- [JKF⁺23] Lucas Joos, Karsten Klein, Maximilian T. Fischer, Frederik L. Dennig, Daniel A. Keim, and Michael Krone. "Exploring Trajectory Data in Augmented Reality: A Comparative Study of Interaction Modalities". In: *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2023, pp. 790–799. DOI: [10.1109/ISMAR59233.2023.00094](https://doi.org/10.1109/ISMAR59233.2023.00094).
- [JMKK19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*. Ed. by Jill Burstein, Christy Doran, and Thamar Solorio. Association for Computational Linguistics, 2019, pp. 4171–4186. DOI: [10.18653/v1/n19-1423](https://doi.org/10.18653/v1/n19-1423).
- [Joh17] Elizabeth E. Joh. "Feeding the Machine: Policing Crime Data & Algorithms". In: *William & Mary Bill of Rights Journal* 26.2 (2017).
- [JP87] David S. Johnson and Henry O. Pollak. "Hypergraph Planarity and the Complexity of Drawing Venn Diagrams". In: *Journal of graph theory* 11.3 (1987), pp. 309–325. DOI: [10.1002/jgt.3190110306](https://doi.org/10.1002/jgt.3190110306).
- [JS18] Jeremy Howard and Sebastian Ruder. "Fine-tuned Language Models for Text Classification". In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (2018), pp. 328–339. DOI: [10.18653/v1/P18-1031](https://doi.org/10.18653/v1/P18-1031).
- [JSS⁺18] Wolfgang Jentner, Rita Sevastjanova, Florian Stoffel, Daniel A. Keim, Jürgen Bernard, and Mennatallah El-Assady. "Minions, Sheep, and Fruits : Metaphorical Narratives to Explain Artificial Intelligence and Build Trust". In: *Workshop on Visualization for AI Explainability at IEEE*. 2018. URL: <http://nbn-resolving.de/urn:nbn:de:bsz:352-2-1q5oi0xg28s59>.
- [JSS⁺22] Wolfgang Jentner, Fabian Sperrle, Daniel Seebacher, Matthias Kraus, Rita Sevastjanova, Maximilian T. Fischer, Udo Schlegel, Dirk Streeb, Matthias Miller, Thilo Spinner, et al. "Visualisierung der COVID-19-Inzidenzen und Behandlungskapazitäten mit CoronaVis". In: *Resilienz und Pandemie*.

Ed. by Andreas Hermann Karsten and Stefan Voßschmidt. Stuttgart: Kohlhammer, 2022, pp. 176–185. ISBN: 978-3-17-039930-3.

- [JWF⁺19] Jianwen Jiang, Yuxuan Wei, Yifan Feng, Jingxuan Cao, and Yue Gao. “Dynamic Hypergraph Neural Networks”. In: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 2019, pp. 2635–2641. ISBN: 978-0-9992411-4-1. DOI: [10.24963/ijcai.2019/366](https://doi.org/10.24963/ijcai.2019/366).
- [JWKN21] Ben Jacobsen, Markus Wallinger, Stephen Kobourov, and Martin Nollenburg. “MetroSets: Visualizing Sets as Metro Maps”. In: *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2021), pp. 1257–1267. DOI: [10.1109/TVCG.2020.3030475](https://doi.org/10.1109/TVCG.2020.3030475).
- [KAF⁺08] Daniel A. Keim, Gennady Andrienko, Jean-Daniel Fekete, Carsten Görg, Jörn Kohlhammer, and Guy Melançon. “Visual Analytics: Definition, Process, and Challenges”. In: *Information Visualization*. Ed. by Andreas Kerren. Springer, 2008, pp. 154–175. ISBN: 978-3-540-70955-8. DOI: [10.1007/978-3-540-70956-5-7](https://doi.org/10.1007/978-3-540-70956-5-7).
- [Kap10] Peter Kapec. “Visualizing software Artifacts Using Hypergraphs”. In: *Proceedings of the 26th Spring Conference on Computer Graphics - SCCG '10*. ACM Press, 2010, p. 27. ISBN: 978-1-4503-0558-7. DOI: [10.1145/1925059.1925067](https://doi.org/10.1145/1925059.1925067).
- [KB15] Diederik P. Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization”. In: *International Conference on Learning Representations (ICLR)* (2015).
- [KBK11] Miloš Krstajić, Enrico Bertini, and Daniel A. Keim. “CloudLines: Compact display of Event Episodes in Multiple Time-series”. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2432–2439. DOI: [10.1109/TVCG.2011.179](https://doi.org/10.1109/TVCG.2011.179).
- [KC04] T. A. Keahey and K. C. Cox. “VIM: A Framework for Intelligence Analysis”. In: *IEEE Symposium on Information Visualization*. IEEE, 2004, p22–p22. ISBN: 0780387793. DOI: [10.1109/INFVIS.2004.72](https://doi.org/10.1109/INFVIS.2004.72).
- [KEL18] Mareile Kaufmann, Simon Egbert, and Matthias Leese. “Predictive Policing and the Politics of Patterns”. In: *The British Journal of Criminology* 59.3 (2018), pp. 674–692. ISSN: 0007-0955. DOI: [10.1093/bjc/azy060](https://doi.org/10.1093/bjc/azy060).
- [KES20] Stefanie Kemme, Iniobong Essien, and Marleen Stelter. “Antimuslimische Einstellungen in der Polizei?: Der Zusammenhang von Kontakthäufigkeit und -qualität mit Vorurteilen und Stereotypen gegenüber Muslimen”. In: *Monatsschrift für Kriminologie und Strafrechtsreform* 103.2 (2020), pp. 129–149. ISSN: 2366-1968. DOI: [10.1515/mks-2020-2048](https://doi.org/10.1515/mks-2020-2048).
- [KFS⁺19] Jay Koven, Cristian Felix, Hossein Siadati, Markus Jakobsson, and Enrico Bertini. “Lessons Learned Developing a Visual Analytics Solution for Investigative Analysis of Scamming Activities”. In: *IEEE Trans. on Visualization and Computer Graphics* 25.1 (2019), pp. 225–234. DOI: [10.1109/TVCG.2018.2865023](https://doi.org/10.1109/TVCG.2018.2865023).

- [KGS09] Youn-ah Kang, Carsten Gorg, and John Stasko. “Evaluating Visual Analytics Systems for Investigative Analysis: Deriving Design Principles From a Case Study”. In: *IEEE Symposium on Visual Analytics Science and Technology*. Ed. by John Stasko and Jarke Jan van Wijk. VAST. Piscataway, NJ, USA: IEEE, 2009, pp. 139–146. ISBN: 978-1-4244-5283-5. DOI: **10.1109/VAST.2009.5333878**.
- [KHT09] Steffen Klamt, Utz-Uwe Haus, and Fabian Theis. “Hypergraphs and Cellular Networks”. In: *PLoS Computational Biology* 5.5 (2009), e1000385. DOI: **10.1371/journal.pcbi.1000385**.
- [Kit17] Rob Kitchin. “Thinking Critically About and Researching Algorithms”. In: *Information, Communication & Society* 20.1 (2017). ISSN: 1369-118X. DOI: **10.1080/1369118X.2016.1154087**.
- [KJ13] Andreas Kerren and Ilir Jusufi. “A Novel Radial Visualization Approach for Undirected Hypergraphs”. In: *EuroVis - Short Papers* (2013), 5 pages. DOI: **10.2312/PE.EuroVisShort.EuroVisShort2013.025-029**.
- [KJZ*21] Omar Shahbaz Khan, Björn Þór Jónsson, Jan Zahálka, Stevan Rudinac, and Marcel Worring. “Impact of Interaction Strategies on User Relevance Feedback”. In: *Proceedings of the 2021 International Conference on Multimedia Retrieval*. Ed. by Wen-Huang Cheng, Mohan Kankanhalli, Meng Wang, Wei-Ta Chu, Jiaying Liu, and Marcel Worring. New York, NY, USA: ACM, 2021, pp. 590–598. ISBN: 9781450384636. DOI: **10.1145/3460426.3463663**.
- [KK15] Kostiantyn Kucher and Andreas Kerren. “Text Visualization Techniques: Taxonomy, Visual Survey, and Community Insights”. In: *2015 IEEE Pacific Visualization Symposium (PacificVis)*. Ed. by Shixia Liu, Gerek Scheuermann, Shigeo Takahashi, and IEEE Pacific Visualization Symposium. Piscataway, NJ: IEEE, 2015, pp. 117–121. ISBN: 978-1-4673-6879-7. DOI: **10.1109/PACIFICVIS.2015.7156366**.
- [KKC15] Natalie Kerracher, Jessie Kennedy, and Kevin Chalmers. “A Task Taxonomy for Temporal Graph Visualisation”. In: *IEEE Transactions on Visualization and Computer Graphics* (2015). DOI: **10.1109/TVCG.2015.2424889**.
- [KKEM10] Daniel A. Keim, Jörn Kohlhammer, Geoffrey P. Ellis, and Florian Mansmann. *Mastering the Information Age Solving Problems with Visual Analytics*. Eurographics Association, 2010. ISBN: 978-3-905673-77-7.
- [KKG20] Hansaka Angel Dias Edirisinghe Kodituwakku, Alex Keller, and Jens Gregor. “InSight2: A Modular Visual Analysis Platform for Network Situational Awareness in Large-Scale Networks”. In: *Electronics* 9.10 (2020), p. 1747. DOI: **10.3390/electronics9101747**.
- [KKL*16] Bum Chul Kwon, Sung-Hee Kim, Sukwon Lee, Jaegul Choo, Jina Huh, and Ji Soo Yi. “VisOHC: Designing Visual Analytics for Online Health Communities”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 71–80. DOI: **10.1109/TVCG.2015.2467555**.

- [KKZE20] Mosab Khayat, Morteza Karimzadeh, Jieqiong Zhao, and David S. Ebert. "VASSL: A Visual Analytics Toolkit for Social Spambot Labeling". In: *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2020), pp. 874–883. doi: [10.1109/TVCG.2019.2934266](https://doi.org/10.1109/TVCG.2019.2934266).
- [KLS07] Bohyoung Kim, Bongshin Lee, and Jinwook Seo. "Visualizing Set Concordance with Permutation Matrices and Fan Diagrams". In: *Interacting with computers* 19.5 (2007), pp. 630–643. ISSN: 0953-5438. doi: [10.1016/j.intcom.2007.05.004](https://doi.org/10.1016/j.intcom.2007.05.004).
- [KM13] Robert Kosara and Jock Mackinlay. "Storytelling: The Next Step for Visualization". In: *Computer* 46.5 (2013), pp. 44–50. ISSN: 0018-9162. doi: [10.1109/MC.2013.36](https://doi.org/10.1109/MC.2013.36).
- [KMN20] Keita Kurita, Paul Michel, and Graham Neubig. "Weight Poisoning Attacks on Pretrained Models". In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Ed. by Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault. Stroudsburg, PA, USA: Association for Computational Linguistics, 2020, pp. 2793–2806. doi: [10.18653/v1/2020.acl-main.249](https://doi.org/10.18653/v1/2020.acl-main.249).
- [KN98] Robert O. Keohane and Joseph S. Nye Jr. "Power and Interdependence in the Information Age". In: *Foreign Affairs* 77 (1998), p. 81.
- [KO11] Fabian Kuhn and Rotem Oshman. "Dynamic Networks: Models and Algorithms". In: *ACM SIGACT News* 42.1 (2011), pp. 82–96. doi: [10.1145/1959045.1959064](https://doi.org/10.1145/1959045.1959064).
- [KP11] Samuel Kaski and Jaakko Peltonen. "Dimensionality Reduction for Data Visualization". In: *IEEE Signal Processing Magazine* 28.2 (2011), pp. 100–104. doi: [10.1109/MSP.2010.940003](https://doi.org/10.1109/MSP.2010.940003).
- [KS11] Youn-ah Kang and John Stasko. "Characterizing the Intelligence Analysis Process: Informing Visual Analytics Design Through a Longitudinal Field Study". In: *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2011, pp. 21–30. ISBN: 978-1-4673-0014-8. doi: [10.1109/VAST.2011.6102438](https://doi.org/10.1109/VAST.2011.6102438).
- [KSS16] Gilad Katz, Eui Chul Richard Shin, and Dawn Song. "ExploreKit: Automatic Feature Generation and Selection". In: *Proceedings of the 16th IEEE International Conference on Data Mining*. Ed. by Francesco Bonchi, Josep Domingo-Ferrer, Ricardo Baeza-Yates, Zhi-Hua Zhou, Xindong Wu, and IEEE International Conference on Data Mining. ICDM. Piscataway, NJ, USA: IEEE, 2016, pp. 979–984. ISBN: 978-1-5090-5473-2. doi: [10.1109/ICDM.2016.0123](https://doi.org/10.1109/ICDM.2016.0123).
- [KV15] James Max Kanter and Kalyan Veeramachaneni. "Deep Feature Synthesis: Towards Automating Data Science Endeavors". In: *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics*. DSAA. IEEE, 2015, pp. 1–10. ISBN: 978-1-4673-8272-4. doi: [10.1109/DSAA.2015.7344858](https://doi.org/10.1109/DSAA.2015.7344858).

- [KW23] Mohan Kankanhalli and Marcel Worring. “Panel: Multimodal Large Foundation Models”. In: *Proceedings of the 31st ACM International Conference on Multimedia*. Ed. by Abdulmotaleb El Saddik, Tao Mei, Rita Cucchiara, Marco Bertini, Diana Patricia Tobon Vallejo, Pradeep K. Atrey, and M. Shamim Hossain. New York, NY, USA: Association for Computing Machinery, 2023, p. 9709. ISBN: 9798400701085. DOI: **10.1145/3581783.3617350**.
- [KY04] Bryan Klimt and Yiming Yang. “The Enron Corpus: A New Dataset for Email Classification Research”. In: *Proceedings of the 15th European Conference on Machine Learning*. ECML’04. Berlin, Heidelberg: Springer-Verlag, 2004, pp. 217–226. ISBN: 3-540-23105-6. DOI: **10.1007/978-3-540-30115-8_22**.
- [LA16] Jeff Larson and Julia Angwin. “Technical Response to Northpointe”. In: *ProPublica* (2016). URL: <https://www.propublica.org/article/technical-response-to-northpointe>.
- [LAH20] Yang Liu, Tim Althoff, and Jeffrey Heer. “Paths Explored, Paths Omitted, Paths Obscured: Decision Points & Selective Reporting in End-to-End Data Analysis”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Ed. by Regina Bernhaupt, Florian ‘Floyd’ Mueller, David Verweij, Josh Andres, Joanna McGrenere, Andy Cockburn, Ignacio Avellino, Alix Goguey, Pernille Bjørn, Shengdong Zhao, et al. New York, NY, USA: ACM, 2020, pp. 1–14. ISBN: 9781450367080. DOI: **10.1145/3313831.3376533**.
- [LDS17] Ji Young Lee, Franck Dernoncourt, and Peter Szolovits. “Transfer Learning for Named-Entity Recognition with Neural Networks”. In: *arXiv preprint arXiv:1705.06273* (2017). DOI: **10.48550/arXiv.1705.06273**.
- [Lea51] Harold J. Leavitt. “Some Effects of Certain Communication Patterns on Group Performance”. In: *The J. of Abnormal and Social Psychology* 46.1 (1951), p. 38.
- [LGYS17] Mingxia Liu, Yue Gao, Pew-Thian Yap, and Dinggang Shen. “Multi-Hypergraph Learning for Incomplete Multimodality Data”. In: *IEEE Journal of Biomedical and Health Informatics* 22.4 (2017), pp. 1197–1208. DOI: **10.1109/JBHI.2017.2732287**.
- [LH19] Ilya Loshchilov and Frank Hutter. “Decoupled Weight Decay Regularization”. In: *International Conference on Learning Representations (ICLR)* (2019).
- [LHKP22] Benjamin Lipp, Paula Helm, Athanasios Karafillidis, and Roser Pujadas. “Theorising Interfaces / Interface Theories”. In: *EASST 2022*. 2022.
- [LI16] Kristian Lum and William Isaac. “To Predict and Serve?” In: *Significance* 13.5 (2016), pp. 14–19. ISSN: 1740-9713. DOI: **10.1111/j.1740-9713.2016.00960.x**.

- [Lip17] Benjamin Lipp. “Analytik des Interfacing. Zur Materialität technologischer Verschaltung in prototypischen Milieus robotisierter Pflege”. In: *1866-2447 10* (2017), pp. 107–129. ISSN: 1866-2447. DOI: **10.6094/behemoth.2017.10.1.948**.
- [Lip22] Walter Lippmann. *Public Opinion*. London: G. Allen & Unwin, 1922.
- [LKT⁺14] Yafeng Lu, Robert Kruger, Dennis Thom, Feng Wang, Steffen Koch, Thomas Ertl, and Ross Maciejewski. “Integrating Predictive Analytics and Social Media”. In: *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2014, pp. 193–202. ISBN: 978-1-4799-6227-3. DOI: **10.1109/VAST.2014.7042495**.
- [LOL⁺18] Bruno Lepri, Nuria Oliver, Emmanuel Letouzé, Alex Pentland, and Patrick Vinck. “Fair, Transparent, and Accountable Algorithmic Decision-making Processes: The Premise, the Proposed Solutions, and the Open Challenges”. In: *Philosophy & Technology* 31.4 (2018), pp. 611–627. ISSN: 2210-5433. DOI: **10.1007/s13347-017-0279-x**.
- [LPP⁺06] Bongshin Lee, Catherine Plaisant, Cynthia Sims Parr, Jean-Daniel Fekete, and Nathalie Henry. “Task Taxonomy for Graph Visualization”. In: *Proceedings of the 2006 AVI workshop on Beyond time and errors novel evaluation methods for information visualization - BELIV '06*. 2006, p. 1. ISBN: 1595935622. DOI: **10.1145/1168149.1168168**.
- [LPT⁺21] Fannie Liu, Chunjong Park, Yu Jiang Tham, Tsung-Yu Tsai, Laura Dabbish, Geoff Kaufman, and Andrés Monroy-Hernández. “Significant Otter: Understanding the Role of Biosignals in Communication”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Ed. by Yoshifumi Kitamura, Aaron Quigley, Katherine Isbister, Takeo Igarashi, Pernille Bjørn, and Steven Drucker. New York, NY, USA: ACM, 5062021, pp. 1–15. ISBN: 9781450380966. DOI: **10.1145/3411764.3445200**.
- [LSB⁺16] Yafeng Lu, Michael Steptoe, Sarah Burke, Hong Wang, Jiun-Yi Tsai, Hasan Davulcu, Douglas Montgomery, Steven R. Corman, and Ross Maciejewski. “Exploring Evolving Media Discourse Through Event Cueing”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 220–229. DOI: **10.1109/TVCG.2015.2467991**.
- [LSB02] Matthew Lombard, Jennifer Snyder-Duch, and Cheryl Campanella Bracken. “Content Analysis in Mass Communication: Assessment and Reporting of Intercoder Reliability”. In: *Human communication research* 28.4 (2002), pp. 587–604.
- [LSHL20] Jing Li, Aixin Sun, Jianglei Han, and Chenliang Li. “A Survey on Deep Learning for Named Entity Recognition”. In: *IEEE Transactions on Knowledge and Data Engineering* 34.1 (2020), pp. 50–70. DOI: **10.1109/TKDE.2020.2981314**.

- [LWY⁺20] Quan Li, Ziming Wu, Lingling Yi, Kristanto Seann, Huamin Qu, and Xiaojuan Ma. “WeSeer: Visual Analysis for Better Information Cascade Prediction of WeChat Articles”. In: *IEEE Transactions on Visualization and Computer Graphics* 26.2 (2020), pp. 1399–1412. DOI: [10.1109/TVCG.2018.2867776](https://doi.org/10.1109/TVCG.2018.2867776).
- [LXLS13] Dong Li, Zhiming Xu, Sheng Li, and Xin Sun. “Link Prediction in Social Networks Based on Hypergraph”. In: *Proceedings of the 22nd International Conference on World Wide Web*. 2013, pp. 41–42. DOI: [10.1145/2487788.2487802](https://doi.org/10.1145/2487788.2487802).
- [LYW⁺16] Shixia Liu, Jialun Yin, Xiting Wang, Weiwei Cui, Kelei Cao, and Jian Pei. “Online Visual Analytics of Text Streams”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.11 (2016), pp. 2451–2466. DOI: [10.1109/TVCG.2015.2509990](https://doi.org/10.1109/TVCG.2015.2509990).
- [LZ15] Qiuju Luo and Dixi Zhong. “Using Social Network Analysis to Explain Communication Characteristics of Travel-related Electronic Word-of-Mouth on Social Networking Sites”. In: *Tourism Management* 46 (2015), pp. 274–282. ISSN: 02615177. DOI: [10.1016/j.tourman.2014.07.007](https://doi.org/10.1016/j.tourman.2014.07.007).
- [Mac15] Adrian Mackenzie. “The Production of Prediction: What Does Machine Learning Want?” In: *European Journal of Cultural Studies* 18.4/5 (2015), pp. 429–445. DOI: [10.1177/1367549415577384](https://doi.org/10.1177/1367549415577384).
- [Mac18] Kevin Macnish. *The Ethics of Surveillance: An Introduction*. The ethics of. London and New York: Routledge Taylor & Francis Group, 2018. ISBN: 978-1138643796.
- [Mäk90] Erkki Mäkinen. “How to Draw a Hypergraph”. In: *International Journal of Computer Mathematics* 34.3-4 (1990), pp. 177–185. DOI: [10.1080/002071690088803875](https://doi.org/10.1080/002071690088803875).
- [MC13] Viktor Mayer-Schönberger and Kenneth Cukier. *Big Data: Die Revolution, die unser Leben verändern wird*. 1. ed. München: Redline-Verlag, 2013. ISBN: 978-3-86881-506-1.
- [MC16] Jacob Metcalf and Kate Crawford. “Where are Human Subjects in Big Data Research? The Emerging Ethics Divide”. In: *Big Data & Society* 3.1 (2016). ISSN: 2053-9517. DOI: [10.1177/2053951716650211](https://doi.org/10.1177/2053951716650211).
- [McL17] Marshall McLuhan. “The Medium Is The Message”. In: *Communication Theory*. Ed. by C. David Mortensen. Routledge, 2017, pp. 390–402. ISBN: 9781315080918. DOI: [10.4324/9781315080918-31](https://doi.org/10.4324/9781315080918-31).
- [McL64] Marshall McLuhan. *Understanding Media: The Extensions of Man*. London: Routledge, 1964. ISBN: 0415253977.
- [MDM⁺15] Sana Malik, Fan Du, Megan Monroe, Eberechukwu Onukwugha, Catherine Plaisant, and Ben Shneiderman. “Cohort Comparison of Event Sequences with Balanced Integration of Visual Analytics and Statistics”. In: *Proceedings of the 20th ACM International Conference on Intelligent User Interfaces*. Ed. by Oliver Brdiczka, Polo Chau, Giuseppe Carenini,

- Shimei Pan, and Per Ola Kristensson. *IUI*. New York, USA: ACM, 2015, pp. 38–49. ISBN: 978-1-4503-3306-1. DOI: **10.1145/2678025.2701407**.
- [Mes09] Gustavo S. Mesch. “Social Context and Communication Channels Choice Among Adolescents”. In: *Computers in Human Behavior* 25.1 (2009), pp. 244–251. ISSN: 07475632. DOI: **10.1016/j.chb.2008.09.007**.
- [MHVB13] Yves-Alexandre de Montjoye, César A. Hidalgo, Michel Verleysen, and Vincent D. Blondel. “Unique in the Crowd: The Privacy Bounds of Human Mobility”. In: *Scientific Reports* 3 (2013), p. 1376. ISSN: 2045-2322.
- [Mit21] Ananda Mitra. *Modern Day Surveillance Ecosystem and Impacts on Privacy*. Hershey: Information Science Reference, 2021. ISBN: 9781799881346.
- [MJ17] Anouk Mols and Susanne Janssen. “Not Interesting Enough to be Followed by the NSA”. In: *Digital Journalism* 5.3 (2017), pp. 277–298. ISSN: 2167-0811. DOI: **10.1080/21670811.2016.1234938**.
- [MKMS20] Angela Mastrianni, Leah Kulp, Emily Mapelli, and Aleksandra Sarcevic. “Understanding Digital Checklist Use Through Team Communication”. In: *Extended Abstracts on Human factors in computing systems. CHI Conference 2020* (2020). DOI: **10.1145/3334480.3382817**.
- [MMKN08] Peter McLachlan, Tamara Munzner, Eleftherios Koutsofios, and Stephen North. “LiveRAC: Interactive Visual Exploration of System Management Time-Series Data”. In: *Proceedings of the 26th CHI Conference on Human Factors in Computing Systems*. Ed. by Mary Czerwinski, Arnie Lund, and Desney Tan. ACM Press, 2008, p. 1483. ISBN: 9781605580111. DOI: **10.1145/1357054.1357286**.
- [MMS⁺21] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. “A Survey on Bias and Fairness in Machine Learning”. In: *ACM Computing Surveys* 54.6 (2021), pp. 1–35. ISSN: 03600300. DOI: **10.1145/3457607**. (Visited on Jan. 12, 2022).
- [Mol19] Christoph Molnar. *Interpretable Machine Learning: A Guide for Making Black Box Models Interpretable*. Morisville, North Carolina: Lulu, 2019. ISBN: 0244768528.
- [Mor34] Jacob L. Moreno. *Who Shall Survive? A New Approach to the Problem of Human Interrelations*. Washington, D.C., 1934. ISBN: 9780266612100.
- [MQB19] Yao Ming, Huamin Qu, and Enrico Bertini. “RuleMatrix: Visualizing and Understanding Classifiers with Rules”. In: *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2019), pp. 342–352. DOI: **10.1109/tvcg.2018.2864812**.
- [MRS⁺13] Wouter Meulemans, Nathalie Henry Riche, Bettina Speckmann, Basak Alper, and Tim Dwyer. “KelpFusion: A Hybrid Set Visualization Technique”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.11 (2013), pp. 1846–1858. DOI: **10.1109/TVCG.2013.76**.
- [MS99] Christopher D. Manning and Hinrich Schütze. *Foundations of Statistical Natural Language Processing*. Cambridge, Mass. and London: MIT Press, 1999. ISBN: 978-0262133609.

- [Müg23] Daniel Mügge. “The Securitization of the EU’s Digital Tech Regulation”. In: *Journal of European Public Policy* (2023), pp. 1–16. ISSN: 1350-1763. DOI: **10.1080/13501763.2023.2171090**.
- [MWBB13] John W. Mohr, Robin Wagner-Pacifici, Ronald L. Breiger, and Petko Bogdanov. “Graphing the Grammar of Motives in National Security Strategies: Cultural Interpretation, Automated Text Analysis and the Drama of Global Politics”. In: *Poetics* 41.6 (2013), pp. 670–700. ISSN: 0304422X. DOI: **10.1016/j.poetic.2013.08.003**.
- [Nat10] National Policing Improvement Agency, ed. *Guidance on the Management of Police Information - Appendix 2*. 2nd. 2010.
- [NGF+20] Minh Hao Nguyen, Jonathan Gruber, Jaelle Fuchs, Will Marler, Amanda Hunsaker, and Eszter Hargittai. “Changes in Digital Communication During the COVID-19 Global Pandemic: Implications for Digital Inequality and Future Research”. In: *Social Media + Society* 6.3 (2020). ISSN: 2056-3051. DOI: **10.1177/2056305120948255**.
- [Nis94] Helen Nissenbaum. “Computing and Accountability”. In: *Communications of the ACM* 37.1 (1994), pp. 72–80. ISSN: 00010782. DOI: **10.1145/175222.175228**.
- [NLH+22] Adela Cañete Nieto, Ruth Ladenstein, Barbara Hero, Sabine Taschner-Mandl, Ulrike Pötschger, Vanessa Düster, Martinez De Las Heras, Blanca, Ana Jiménez Pastor, Eva Bozsaky, Maximilian T. Fischer, et al. “PRIMAGE - An Artificial Intelligence-based Clinical Decision Support System for Optimized Cancer Diagnosis and Risk Assessment - A Progress Update”. In: *SIOPEX AGM*. Athens, 2022.
- [NNVW15] Hong-Wei Ng, Viet Dung Nguyen, Vassilios Vonikakis, and Stefan Winkler. “Deep Learning for Emotion Recognition on Small Datasets using Transfer Learning”. In: *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*. Ed. by Zhengyou Zhang, Phil Cohen, Dan Bohus, Radu Horaud, and Helen Meng. New York: ACM, 2015, pp. 443–449. ISBN: 9781450339124. DOI: **10.1145/2818346.2830593**.
- [Nob18] Safiya Umoja Noble. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018. ISBN: 9781479849949.
- [NS07] David Nadeau and Satoshi Sekine. “A Survey of Named Entity Recognition and Classification”. In: *Linguisticae Investigationes* 30.1 (2007), pp. 3–26. DOI: **10.1075/li.30.1.03nad**.
- [Nui20] Nuix Pty Ltd. *Nuix Discover and Nuix Investigate*. 2020. URL: **<https://www.nuix.com/products>**.
- [NXWW16] Phong H. Nguyen, Kai Xu, Rick Walker, and B. WilliamL Wong. “TimeSets: Timeline Visualization With Set Relations”. In: *Information Visualization* 15.3 (2016), pp. 253–269. ISSN: 1473-8724. DOI: **10.1177/1473871615605347**.

- [OLM17] Xavier Ouvrard, Jean-Marie Le Goff, and Stéphane Marchand-Maillet. “Networks of Collaborations: Hypergraph Modeling and Visualisation”. In: *CoRR abs/1707.00115* (2017).
- [ONe16] Cathy O’Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. 1st edition. Penguin Books UK, 2016. ISBN: 978-0-14-198541-1.
- [OO19] Ijeoma Onyeator and Ngozi Okpara. “Human Communication in a Digital Age: Perspectives on Interpersonal communication in the family”. In: *New Media and Mass Communication* 78 (2019), pp. 35–45. DOI: **10.7176/NMMC/78-06**.
- [Ope23] OpenAI. *GPT-4 Technical Report*. 2023. DOI: **10.48550/arXiv.2303.08774**.
- [OS23] Anke Sophia Obendiek and Timo Seidl. “The (False) Promise of Solutionism: Ideational Business Power and the Construction of Epistemic Authority in Digital Security Governance”. In: *Journal of European Public Policy* (2023), pp. 1–25. ISSN: 1350-1763. DOI: **10.1080/13501763.2023.2172060**.
- [OSH+07] Jukka-Pekka Onnela, Jari Saramäki, Jörkki Hyvönen, Gábor Szabó, M. Argollo de Menezes, Kimmo Kaski, Albert-László Barabási, and János Kertész. “Analysis of a Large-Scale Weighted Network of One-to-One Human Communication”. In: *New Journal of Physics* 9.6 (2007), p. 179. DOI: **10.1088/1367-2630/9/6/179**.
- [Pal20] Palantir Technologies, Inc. *Gotham*. 2020. URL: <https://www.palantir.com/palantir-gotham/>.
- [Pao02] Letizia Paoli. “The Paradoxes of Organized Crime”. In: *Crime, Law and Social Change* 37.1 (2002), pp. 51–97. ISSN: 1573-0751. DOI: **10.1023/A:1013355122531**.
- [Pas15] Frank Pasquale. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press, 2015. ISBN: 978-0-674-36827-9.
- [PBL17] Ashwin Paranjape, Austin R. Benson, and Jure Leskovec. “Motifs in Temporal Networks”. In: *Proceedings of the International Conference on Web and Data Mining - WSDM*. ACM Press, 2017, pp. 601–610. ISBN: 978-1-4503-4675-7. DOI: **10.1145/3018661.3018731**.
- [PBN17] John J. Palowitch, Shankar Bhamidi, and Andrew B. Nobel. *Significance-based community detection in weighted networks*. 2017. URL: <http://arxiv.org/pdf/1601.05630v4>.
- [Pea11] Judy C. Pearson. *Human Communication*. 4th. McGraw-Hill, 2011. ISBN: 978-0073406800.
- [PGU+13] Adam Perer, Ido Guy, Erel Uziel, Inbal Ronen, and Michal Jacovi. “The Longitudinal Use of SaNDVis: Visual social Network Analytics in the Enterprise”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.7 (2013), pp. 1095–1108. DOI: **10.1109/TVCG.2012.322**.

- [PL08] Bo Pang and Lillian Lee. "Opinion Mining and Sentiment Analysis". In: *Foundations and Trends in Inf. Retr.* 2.1-2 (2008), pp. 1–135. DOI: [10.1561/1/1500000011](https://doi.org/10.1561/1/1500000011).
- [Pla11] Plato. *Gorgias*. Stuttgart: Reclam, 2011. ISBN: 978-3150188965.
- [PP05] E. Ashby Plant and B. Michelle Peruche. "The Consequences of Race for Police Officers' Responses to Criminal Suspects". In: *Psychological Science* 16.3 (2005), pp. 180–183. ISSN: 0956-7976. DOI: [10.1111/j.0956-7976.2005.00800.x](https://doi.org/10.1111/j.0956-7976.2005.00800.x).
- [Prž11] Nataša Pržulj. "Protein-Protein Interactions: Making Sense of Networks Via Graph-Theoretic Modeling". In: *Bioessays* 33.2 (2011), pp. 115–123. DOI: [10.1002/bies.201000044](https://doi.org/10.1002/bies.201000044).
- [PT11] Jesse Paquette and Taku Tokuyasu. "Hypergraph Visualization and Enrichment Statistics: How the EGAN Paradigm Facilitates Organic Discovery from Big Data". In: *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)* 7865 (2011). DOI: [10.1117/12.890220](https://doi.org/10.1117/12.890220).
- [PUL22] Andrei Paleyes, Raoul-Gabriel Urma, and Neil D. Lawrence. "Challenges in Deploying Machine Learning: A Survey of Case Studies". In: *ACM Computing Surveys* 55.6 (2022), pp. 1–29. ISSN: 03600300. DOI: [10.1145/3533378](https://doi.org/10.1145/3533378).
- [PVW20] Roser Pujadas, Erika Valderrama, and Will Venters. "Interfaces and the Dynamics of Digital Ecosystems: A Study of the Online Travel Ecosystem". In: *International Conference on Information Systems (ICIS)* (2020).
- [Rei10] Florian Reitz. "A Framework for an Ego-centered and Time-aware Visualization of Relations in Arbitrary Data Repositories". In: *arXiv abs/1009.5183* (2010). DOI: [10.48550/arXiv.1009.5183](https://doi.org/10.48550/arXiv.1009.5183).
- [Rib23] Jordi Ribas. *Building the New Bing*. 2023. URL: https://www.linkedin.com/pulse/building-new-bing-jordi-ribas/?src=aff-ref&trk=aff-ir_progid.8005_partid.10078_sid._adid.449670.
- [Rig19] Christopher Rigano. "Using Artificial Intelligence to Address Criminal Justice Needs". In: *National Institute of Justice Journal* 280.1-10 (2019), p. 17.
- [RK80] Everett M. Rogers and D. Lawrence Kincaid. *Communication Networks: Toward a New Paradigm for Research*. Macmillan USA, 1980. ISBN: 978-0029267400.
- [RKX+22] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. *Robust Speech Recognition via Large-Scale Weak Supervision*. 2022. DOI: [10.48550/arXiv.2212.04356](https://doi.org/10.48550/arXiv.2212.04356).
- [RR09] Lev Ratinov and Dan Roth. "Design Challenges and Misconceptions in Named Entity Recognition". In: *Proceedings of the Thirteenth Conference on Computational Natural Language Learning*. CoNLL '09. USA: Association for Computational Linguistics, 2009, pp. 147–155. ISBN: 9781932432299.

- [ŘS10] Radim Řehůřek and Petr Sojka. “Software Framework for Topic Modelling with Large Corpora”. In: *Proceedings of LREC 2010 workshop: New Challenges for NLP Frameworks*. Valletta, Malta, 2010, pp. 46–50. ISBN: 2-9517408-6-7. URL: <http://www.fi.muni.cz/usr/sojka/presentations/lrec2010-poster-rehurek-sojka.pdf>.
- [RTK*14] Anna Ritz, Allison N. Tegge, Hyunju Kim, Christopher L. Poirel, and T. M. Murali. “Signaling Hypergraphs”. In: *Trends in Biotechnology* 32.7 (2014), pp. 356–362. doi: [10.1016/j.tibtech.2014.04.007](https://doi.org/10.1016/j.tibtech.2014.04.007).
- [Rud19] Cynthia Rudin. “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead”. In: *Nature Machine Intelligence* 1.5 (2019), pp. 206–215. ISSN: 2522-5839. doi: [10.1038/s42256-019-0048-x](https://doi.org/10.1038/s42256-019-0048-x).
- [SAKW19] Dirk Streeb, Devanshu Arya, Daniel A. Keim, and Marcel Worring. “Visual Analytics Framework for the Assessment of Temporal Hypergraph Prediction Models”. In: *Proceedings of the Set Visual Analytics Workshop at IEEE VIS 2019*. 2019.
- [SBDK05] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph ‘Jofish’ Kaye. “Reflective Design”. In: *Proceedings of the 4th Decennial Conference on Critical Computing Between Sense and Sensibility - CC ’05*. Aarhus, Denmark: ACM Press, 2005, p. 49. ISBN: 978-1-59593-203-7. doi: [10.1145/1094562.1094569](https://doi.org/10.1145/1094562.1094569).
- [SC21] Klaus Schwarz and Reiner Creutzburg. “Design of Professional Laboratory Exercises for Effective State-of-the-Art OSINT Investigation Tools - Part 3: Maltego”. In: *Electronic Imaging* 33.3 (2021), pp. 45-1-45-23. doi: [10.2352/ISSN.2470-1173.2021.3.MOBMU-045](https://doi.org/10.2352/ISSN.2470-1173.2021.3.MOBMU-045).
- [SCE22] Fabian Sperrle, Davide Ceneda, and Mennatallah El-Assady. “Lotse: A Practical Framework for Guidance in Visual Analytics”. In: *IEEE Transactions on Visualization and Computer Graphics* PP (2022). doi: [10.1109/TVCG.2022.3209393](https://doi.org/10.1109/TVCG.2022.3209393).
- [Sch05] Jean Scholtz. “Metrics for Evaluation of Software Technology to Support Intelligence Analysis”. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 49.10 (2005), pp. 918–921. ISSN: 2169-5067. doi: [10.1177/154193120504901010](https://doi.org/10.1177/154193120504901010).
- [Sch22] Dominik Schwimmbeck. *HuggingFace: Domischwimmbeck/Bert-base-german-cased-fine-tuned-ner*. 2022. URL: <https://huggingface.co/domischwimmbeck/bert-base-german-cased-fine-tuned-ner>.
- [Sch81] Friedemann Schulz von Thun. *Miteinander reden: Störungen und Klärungen - Psychologie der zwischenmenschlichen Kommunikation*. Rowohlt, 1981. ISBN: 3499174898.
- [Sco09] Carlos Alberto Scolari. “Mapping Conversations About New Media: The Theoretical Field of Digital Communication”. In: *New Media & Society* 11.6 (2009), pp. 943–964. ISSN: 1461-4448. doi: [10.1177/1461444809336513](https://doi.org/10.1177/1461444809336513).

- [Sco17] John Scott. *Social Network Analysis*. 4th. Los Angeles: SAGE, 2017. ISBN: 9781473952126.
- [Sco22] Mark Scott. "How Ukraine used Russia's digital playbook against the Kremlin". In: *POLITICO* (2022). URL: <https://www.politico.eu/article/ukraine-russia-digital-playbook-war/>.
- [SD60] I. Richard Savage and Karl W. Deutsch. "A Statistical Model of the Gross Analysis of Transaction Flows". In: *Econometrica* 28.3 (1960), pp. 551–572. ISSN: 00129682. URL: <http://www.jstor.org/stable/1910131>.
- [Sel17] Andrew D. Selbst. "Disparate Impact in Big Data Policing". In: *Georgia Law Review* 109 52.1 (2017), p. 3373.
- [SFE15] Feng Shi, Jacob G. Foster, and James A. Evans. "Weaving the Fabric of Science: Dynamic Network Models of Science's Unfolding Structure". In: *Social Networks* 43 (2015), pp. 73–85. ISSN: 03788733. DOI: [10.1016/j.socnet.2015.02.006](https://doi.org/10.1016/j.socnet.2015.02.006).
- [SFS⁺19] Daniel Seebacher, Maximilian T. Fischer, Rita Sevastjanova, Daniel A. Keim, and Mennatallah El-Assady. "Visual Analytics of Conversational Dynamics". In: *EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by Tatiana von Landesberger and Cagatay Turkey. EuroVA. Porto, Portugal: The Eurographics Association, 2019. ISBN: 978-3-03868-087-1. DOI: [10.2312/eurova.20191130](https://doi.org/10.2312/eurova.20191130).
- [SGLS07] John Stasko, Carsten Gorg, Zhicheng Liu, and Kanupriya Singhal. "Jigsaw: Supporting Investigative Analysis through Interactive Visualization". In: *2007 IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 2007, pp. 131–138. ISBN: 978-1-4244-1659-2. DOI: [10.1109/VAST.2007.4389006](https://doi.org/10.1109/VAST.2007.4389006).
- [SGS18] A. Sarikaya, M. Gleicher, and D. A. Szafrir. "Design Factors for Summary Visualization in Visual Analytics". In: *Computer Graphics Forum* 37.3 (2018), pp. 145–156. ISSN: 01677055. DOI: [10.1111/cgf.13408](https://doi.org/10.1111/cgf.13408).
- [SHMK14] Michael Sailer, Jan Hense, J. Mandl, and Markus Klevers. "Psychological Perspectives on Motivation Through Gamification". In: *Interaction Design and Architecture Journal* 19 (2014), pp. 28–37.
- [Shn20] Ben Shneiderman. "Bridging the Gap Between Ethics and Practice". In: *ACM Transactions on Interactive Intelligent Systems* 10.4 (2020), pp. 1–31. ISSN: 2160-6455. DOI: [10.1145/3419764](https://doi.org/10.1145/3419764).
- [Sim08] Georg Simmel. *Soziologie: Untersuchungen über die Formen der Vergesellschaftung*. Leipzig: Duncker & Humblot, 1908.
- [Sin09] Simon Sinek. *Start With Why: How Great Leaders Inspire Everyone to Take Action*. Penguin, 2009. ISBN: 978-1-59184-644-4.
- [SJB⁺17] Florian Stoffel, Wolfgang Jentner, Michael Behrisch, Johannes Fuchs, and Daniel A. Keim. "Interactive Ambiguity Resolution of Named Entities in Fictional Literature". In: *Computer Graphics Forum* 36.3 (2017), pp. 189–200. ISSN: 01677055. DOI: [10.1111/cgf.13179](https://doi.org/10.1111/cgf.13179).

- [SJB*20] Fabian Sperrle, Astrik Jeitler, Jürgen Bernard, Daniel A. Keim, and Mennatallah El-Assady. “Learning and Teaching in Co-Adaptive Guidance for Mixed-Initiative Visual Analytics”. In: *EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by Cagatay Turkay and Katerina Vrotsou. 2020. DOI: **10.2312/eurova.20201088**.
- [SJB*21] Fabian Sperrle, Astrik Jeitler, Jürgen Bernard, Daniel A. Keim, and Mennatallah El-Assady. “Co-Adaptive Visual Data Analysis and Guidance Processes”. In: *Computers & Graphics* 100 (2021), pp. 93–105. ISSN: 00978493. DOI: **10.1016/j.cag.2021.06.016**.
- [SK16] Tim Salimans and Durk P. Kingma. “Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 901–909. DOI: **10.5555/3157096.3157197**.
- [SLN05] Purvi Saraiya, Peter Lee, and Chris North. “Visualization of Graphs with Associated Timeseries Data”. In: *Proceedings of the IEEE Symposium on Information Visualization*. Ed. by John Stasko and Matthew O. Ward. InfoVis. IEEE, 2005, pp. 225–232. ISBN: 0-7803-9464-x. DOI: **10.1109/INFVIS.2005.1532151**.
- [SMC20] P. G. Shynu, H. Md. Shayan., and Chiranji Lal Chowdhary. “A Fuzzy based Data Perturbation Technique for Privacy Preserved Data Mining”. In: *2020 International Conference on Emerging Trends in Information Technology and Engineering (ic-ETITE)*. IEEE, 2020, pp. 1–4. ISBN: 978-1-7281-4142-8. DOI: **10.1109/ic-ETITE47903.2020.244**.
- [Spa91] Malcolm K. Sparrow. “The Application of Network Analysis to Criminal Intelligence: An Assessment of the Prospects”. In: *Social Networks* 13.3 (1991), pp. 251–274. ISSN: 03788733. DOI: **10.1016/0378-8733(91)90008-H**.
- [Spi17] Matthias Spielkamp. “Inspecting Algorithms for Bias”. In: *MIT Technology Review* (2017). URL: <https://www.technologyreview.com/2017/06/12/105804/inspecting-algorithms-for-bias/>.
- [SRG*19] C. Stoiber, A. Rind, F. Grassinger, R. Gutounig, E. Goldgruber, M. Sedlmair, Š. Emrich, and W. Aigner. “Netflower: Dynamic Network Visualization for Data Journalists”. In: *Computer Graphics Forum* 38.3 (2019), pp. 699–711. ISSN: 01677055. DOI: **10.1111/cgf.13721**.
- [SRKS16] Nivethitha Somu, M. R. Gauthama Raman, Kannan Kirthivasan, and V. S. Shankar Sriram. “Hypergraph Based Feature Selection Technique for Medical Diagnosis”. In: *Journal of Medical Systems* 40 (2016), p. 239. ISSN: 0148-5598. DOI: **10.1007/s10916-016-0600-8**.
- [SSG12] Hadi Shiravi, Ali Shiravi, and Ali A. Ghorbani. “A Survey of Visualization Systems for Network Security”. In: *Transactions on Visualization and Computer Graphics* 18.8 (2012), pp. 1313–1329. ISSN: 1077-2626. DOI: **10.1109/TVCG.2011.144**.

- [SSK⁺16] Dominik Sacha, Hansi Senaratne, Bum Chul Kwon, Geoffrey Ellis, and Daniel A. Keim. "The Role of Uncertainty, Awareness, and Trust in Visual Analytics". In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 240–249. DOI: [10.1109/TVCG.2015.2467591](https://doi.org/10.1109/TVCG.2015.2467591).
- [SSS⁺14] Dominik Sacha, Andreas Stoffel, Florian Stoffel, Bum Chul Kwon, Geoffrey P. Ellis, and Daniel A. Keim. "Knowledge Generation Model for Visual Analytics". In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1604–1613. DOI: [10.1109/TVCG.2014.2346481](https://doi.org/10.1109/TVCG.2014.2346481).
- [SSSE20] Thilo Spinner, Udo Schlegel, Hanna Schafer, and Mennatallah El-Assady. "explAIner: A Visual Analytics Framework for Interactive and Explainable Machine Learning". In: *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2020), pp. 1064–1074. DOI: [10.1109/TVCG.2019.2934629](https://doi.org/10.1109/TVCG.2019.2934629).
- [Str19] Jonathan Stray. "Making Artificial Intelligence Work for Investigative Journalism". In: *Digital Journalism* 7.8 (2019), pp. 1076–1097. ISSN: 2167-0811. DOI: [10.1080/21670811.2019.1630289](https://doi.org/10.1080/21670811.2019.1630289).
- [STT⁺20] Raphael Scheible, Fabian Thomczyk, Patric Tippmann, Victor Jaravine, and Martin Boeker. "GottBERT: A Pure German Language Model". In: *arXiv preprint arXiv:2012.02110* (2020).
- [Suc07] Lucille Alice Suchman. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd. Cambridge: Cambridge University Press, 2007. ISBN: 9780521675888.
- [SV19] Khan Sameera and Pinki Vishwakarma. "Cybercrime: To Detect Suspected User's Chat Using Text Mining". In: *Information and Communication Technology for Intelligent Systems*. Ed. by Suresh Chandra Satapathy and Amit Joshi. Vol. 106. Smart Innovation, Systems and Technologies. Singapore: Springer Singapore, 2019, pp. 381–390. ISBN: 978-981-13-1741-5. DOI: [10.1007/978-981-13-1742-2-37](https://doi.org/10.1007/978-981-13-1742-2-37).
- [SW49] Claude Elwood Shannon and W. Weaver. *Mathematical Theory of Communication*. Univ Illinois Press, 1949. ISBN: 9780252725487.
- [SZS⁺17] Dominik Sacha, Leishi Zhang, Michael Sedlmair, John Aldo Lee, Jaakko Peltonen, Daniel Weiskopf, Stephen C. North, and Daniel A. Keim. "Visual Interaction with Dimensionality Reduction: A Structured Literature Analysis". In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2017), pp. 241–250. DOI: [10.1109/TVCG.2016.2598495](https://doi.org/10.1109/TVCG.2016.2598495).
- [THW⁺21] Nguyen Anh Tu, Thien Huynh-The, Kok-Seng Wong, M. Fatih Demirci, and Young-Koo Lee. "Toward Efficient and Intelligent Video Analytics with Visual Privacy Protection for Large-Scale Surveillance". In: *The Journal of Supercomputing* 77.12 (2021), pp. 14374–14404. ISSN: 0920-8542. DOI: [10.1007/s11227-021-03865-7](https://doi.org/10.1007/s11227-021-03865-7).
- [TLV⁺19] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. "HuggingFace's Transformers: State-of-the-art Natural Language Processing". In: *CoRR abs/1910.03771* (2019).

- [Tom08] Michael Tomasello. *Origins of Human Communication*. Vol. 2008. The Jean Nicod lectures. Cambridge, Mass. and London: MIT, 2008. ISBN: 978-0-262-20177-3.
- [Tri08] Matthias Trier. “Research Note –Towards Dynamic Visualization for Understanding Evolution of Digital Communication Networks”. In: *Information Systems Research* 19.3 (2008), pp. 335–350. ISSN: 1047-7047. DOI: [10.1287/isre.1080.0191](https://doi.org/10.1287/isre.1080.0191).
- [TWW⁺22] Tan Tang, Yanhong Wu, Yingcai Wu, Lingyun Yu, and Yuhong Li. “Video-Moderator: A Risk-aware Framework for Multimodal Video Moderation in E-Commerce”. In: *IEEE Transactions on Visualization and Computer Graphics* 28.1 (2022), pp. 846–856. DOI: [10.1109/TVCG.2021.3114781](https://doi.org/10.1109/TVCG.2021.3114781).
- [UNO11] UNODC. *Criminal Intelligence: Manual for Analysts*. Vienna, Austria, 2011. URL: https://www.unodc.org/documents/organized-crime/Law-Enforcement/Criminal_Intelligence_for_Analysts.pdf.
- [van05] J. J. van Wijk. “The Value of Visualization”. In: *VIS 05. IEEE Visualization, 2005*. IEEE, 2005, pp. 79–86. ISBN: 0-7803-9462-3. DOI: [10.1109/visual.2005.1532781](https://doi.org/10.1109/visual.2005.1532781).
- [VBF17] Paola Valdivia, Paolo Buono, and Jean-Daniel Fekete. “Hypenet: Visualizing Dynamic Hypergraphs”. In: *EuroVis 2017 - Posters*. The Eurographics Association, 2017. DOI: [10.2312/eurp.20171162](https://doi.org/10.2312/eurp.20171162).
- [VBP⁺18] Paola Valdivia, Paolo Buono, Catherine Plaisant, Nicole Dufournaud, and Jean-Daniel Fekete. “Using Dynamic Hypergraphs to Reveal the Evolution of the Business Network of a 17th Century French Woman Merchant”. In: *VIS 2018 - 3rd Workshop on Visualization for the Digital Humanities*. Berlin, Germany, 2018. URL: <https://hal.inria.fr/hal-01855675>.
- [VBP⁺19] Paola Valdivia, Paolo Buono, Catherine Plaisant, Nicole Dufournaud, and Jean-Daniel Fekete. “Analyzing Dynamic Hypergraphs with Parallel Aggregated Ordered Hypergraph Visualization”. In: *IEEE Transactions on Visualization and Computer Graphics* (2019). DOI: [10.1109/TVCG.2019.2933196](https://doi.org/10.1109/TVCG.2019.2933196).
- [VBW17] Corinna Vehlow, Fabian Beck, and Daniel Weiskopf. “Visualizing Group Structures in Graphs: A Survey”. In: *Computer Graphics Forum* 36.6 (2017), pp. 201–225. ISSN: 01677055. DOI: [10.1111/cgf.12872](https://doi.org/10.1111/cgf.12872).
- [VGD06] Fernanda B. Viégas, Scott Golder, and Judith Donath. “Visualizing email content”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Ed. by Rebecca Grinter, Thomas Rodden, Paul Aoki, Ed Cutrell, Robin Jeffries, and Gary Olson. ACM, 2006, pp. 979–988. ISBN: 1595933727. DOI: [10.1145/1124772.1124919](https://doi.org/10.1145/1124772.1124919).

- [vP18] Wouter van Atteveldt and Tai-Quan Peng. “When Communication Meets Computation: Opportunities, Challenges, and Pitfalls in Computational Communication Science”. In: *Communication Methods and Measures* 12.2-3 (2018), pp. 81–92. ISSN: 1931-2458. DOI: **10.1080/19312458.2018.1458084**.
- [VSP⁺17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. “Attention Is All You Need”. In: *Advances in neural information processing systems* 30 (2017).
- [vW14] Stef van den Elzen and Jarke J. van Wijk. “Multivariate Network Exploration and Presentation: From Detail to Overview via Selections and Aggregations”. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2310–2319. DOI: **10.1109/TVCG.2014.2346441**.
- [WBJ74] Paul Watzlawick, Janet H. Beavin, and D. D. J. Jackson. *Menschliche Kommunikation: Formen, Störungen, Paradoxien*. 4th ed. Hans Huber, 1974. ISBN: 3456306105.
- [WCG⁺16] Yingcai Wu, Nan Cao, David Gotz, Yap-Peng Tan, and Daniel A. Keim. “A Survey on Visual Analytics of Social Media Data”. In: *IEEE Transactions on Multimedia* 18.11 (2016), pp. 2135–2148. ISSN: 1520-9210. DOI: **10.1109/TMM.2016.2614220**.
- [WLY⁺14] Yingcai Wu, Shixia Liu, Kai Yan, Mengchen Liu, and Fangzhao Wu. “OpinionFlow: Visual Analysis of Opinion Diffusion on Social Media”. In: *IEEE Transactions on Visualization and Computer Graphics* (2014). DOI: **10.1109/TVCG.2014.2346920**.
- [WQ20] Aoyu Wu and Huamin Qu. “Multimodal Analysis of Video Collections: Visual Exploration of Presentation Techniques in TED Talks”. In: *IEEE Transactions on Visualization and Computer Graphics* 26.7 (2020), pp. 2429–2442. DOI: **10.1109/TVCG.2018.2889081**.
- [WSE19] Emily Wall, John Stasko, and Alex Endert. “Toward a Design Space for Mitigating Cognitive Bias in Vis”. In: *2019 IEEE Visualization Conference (VIS)*. IEEE, 2019, pp. 111–115. ISBN: 978-1-7281-4941-7. DOI: **10.1109/VISUAL.2019.8933611**.
- [WXLW07] Yong-Jie Wang, Ming Xian, Jin Liu, and Guo-yu Wang. “Study of Network Security Evaluation Based on Attack Graph Model”. In: *Journal of China Institute of Communications* 28.3 (2007), p. 29.
- [WYB18] Gregor Wiedemann, Seid Muhie Yimam, and Chris Biemann. “New/s/leak 2.0 – Multilingual Information Extraction and Visualization for Investigative Journalism”. In: *Social Informatics*. Ed. by Steffen Staab, Olessia Koltsova, and Dmitry I. Ignatov. Vol. 11186. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2018, pp. 313–322. ISBN: 978-3-030-01158-1. DOI: **10.1007/978-3-030-01159-8_30**.

- [XDC⁺13] Panpan Xu, Fan Du, Nan Cao, Conglei Shi, Hong Zhou, and Huamin Qu. “Visual Analysis of Set Relations in a Graph”. In: *Computer Graphics Forum* 32.3pt1 (2013), pp. 61–70. ISSN: 01677055. DOI: **10.1111/cgf.12093**.
- [XKS13] Jierui Xie, Stephen Kelley, and Boleslaw K. Szymanski. “Overlapping Community Detection in Networks: the State of the Art and Comparative Study”. In: *ACM Computing Surveys* 45.4 (2013), pp. 1–35. ISSN: 03600300. DOI: **10.1145/2501654.2501657**.
- [XLH⁺11] Jing Xia, Zhen Liu, Yueqi Hu, Wei Chen, and Dichao Peng. “Hypergraph-based bone dataset visualization”. In: *Journal of Computer Aided Design & Computer Graphics* 23.12 (2011), pp. 2040–2045.
- [XSL11] Jierui Xie, Boleslaw K. Szymanski, and Xiaoming Liu. “SLPA: Uncovering Overlapping Communities in Social Networks via A Speaker-listener Interaction Dynamic Process”. In: *Proceedings of the International Conference on Data Mining Workshops - ICDMW*. IEEE, 2011, pp. 344–349. ISBN: 978-1-4673-0005-6. DOI: **10.1109/ICDMW.2011.154**.
- [YEL10] Ji Soo Yi, Niklas Elmqvist, and Seungyoon Lee. “TimeMatrix: Analyzing Temporal Social Networks Using Interactive Matrix-Based Visualizations”. In: *International Journal of Human-Computer Interaction* 26.11-12 (2010), pp. 1031–1051. DOI: **10.1080/10447318.2010.516722**.
- [YKSJ07] Ji Soo Yi, Youn Ah Kang, John Stasko, and Julie Jacko. “Toward a Deeper Understanding of the Role of Interaction in Information Visualization”. In: *Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1224–1231. ISSN: 1077-2626. DOI: **10.1109/TVCG.2007.70515**.
- [YNY⁺19] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar. “HyperGCN: A New Method For Training Graph Convolutional Networks on Hypergraphs”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 1509–1520. DOI: **10.5555/3454287.3454422**.
- [YPO4] Byungun Yoon and Yongtae Park. “A Text-Mining-Based Patent Network: Analytical Tool for High-Technology Trend”. In: *Journal of High Technology Management Research* 15.1 (2004), pp. 37–50. DOI: **10.1016/j.hitech.2003.09.003**.
- [Zar16] Tal Zarsky. “The Trouble with Algorithmic Decisions: An Analytic Road Map to Examine Efficiency and Fairness in Automated and Opaque Decision Making”. In: *Science, Technology & Human Values* 41.1 (2016), pp. 118–132. ISSN: 0162-2439. DOI: **10.1177/0162243915605575**.
- [ZCCB12] Jian Zhao, F. Chevalier, C. Collins, and R. Balakrishnan. “Facilitating Discourse Analysis with Interactive Visualization”. In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2639–2648. DOI: **10.1109/TVCG.2012.226**.

- [ZHS07] Dengyong Zhou, Jiayuan Huang, and Bernhard Schölkopf. “Learning with Hypergraphs: Clustering, Classification, and Embedding”. In: *Advances in Neural Information Processing Systems*. 2007, pp. 1601–1608. DOI: **10.5555/2976456.2976657**.
- [ZLD⁺15] Jian Zhao, Zhicheng Liu, Mira Dontcheva, Aaron Hertzmann, and Alan Wilson. “MatrixWave”. In: *Proceedings of the 33rd CHI Conference on Human Factors in Computing Systems*. Ed. by Bo Begole, Jinwoo Kim, Kori Inkpen, and Woontack Woo. ACM, 2015, pp. 259–268. ISBN: 978-1-4503-3145-6. DOI: **10.1145/2702123.2702419**.
- [ZLW22] Alexandra Zytek, Dongyu Liu, Rhema Vaithianathan, and Kalyan Veeramachaneni. “Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making”. In: *IEEE Transactions on Visualization and Computer Graphics* 28.1 (2022), pp. 1161–1171. DOI: **10.1109/TVCG.2021.3114864**.
- [ZW14] Jan Zahalka and Marcel Worring. “Towards Interactive, Intelligent, and Integrated Multimedia Analytics”. In: *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 2014, pp. 3–12. ISBN: 978-1-4799-6227-3. DOI: **10.1109/VAST.2014.7042476**.
- [ZWW⁺20] Haipeng Zeng, Xingbo Wang, Aoyu Wu, Yong Wang, Quan Li, Alex Endert, and Huamin Qu. “EmoCo: Visual Analysis of Emotion Coherence in Presentation Videos”. In: *IEEE Transactions on Visualization and Computer Graphics* 26.1 (2020), pp. 927–937. DOI: **10.1109/TVCG.2019.2934656**.