Computers & Graphics (2021)

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Co-Adaptive Visual Data Analysis and Guidance Processes

Fabian Sperrle[a,*], Astrik Jeitler[a], Jürgen Bernard[b], Daniel Keim[a], Mennatallah El-Assady[a]

[a]*University of Konstanz, Germany*
[b]*University of Zurich, Switzerland*

ARTICLE INFO

Article history:
Received September 21, 2021

Keywords: Co-Adaptive Analysis Process, Guidance, Visual Analytics

ABSTRACT

Mixed-initiative visual data analysis processes are characterized by the co-adaptation of users and systems over time. As the analysis progresses, both actors – users and systems – gather information, update their analysis behavior, and work on different tasks towards their respective goals. In this paper, we contribute a multigranular model of co-adaptive visual analysis that is centered around incremental learning goals derived from a hierarchical taxonomy of learning goals from pedagogy. Our model captures how both actors adapt their data-, task-, and user/system-models over time. We characterize interaction patterns in terms of the dynamics of learning and teaching that drive adaptation. To demonstrate our model's applicability, we outline aspects of co-adaptation in related models of visual analytics and highlight co-adaptation in existing applications. We further postulate a set of expectations towards adaptation in mixed-initiative processes and identify open research questions and opportunities for future work in co-adaptation.

1. Introduction

Mixed-initiative interaction [1] is at the core of visual analytics, a field of research that aims to combine human intuition and domain knowledge with automated data analysis and visualization. Early visual analytics approaches often relied on intuitive visual representations of data and patterns to support users in their analysis [2]. More recently, systems are taking on an increasingly *active* role in the mixed-initiative process [1]. This has renewed interest in active guidance, "a computer-assisted process that aims to actively resolve a knowledge-gap encountered by users during an interactive visual analytics session" [3]. Other definitions have called for guidance to provide "*just-in-time*" facilitation [4] and specify that guidance should be contextualized and able to adapt to different scenarios dynamically [4].

While the current definition of guidance captures the mixed-initiative nature of the process, it does not shed light on how users and systems adapt over time. Hence, there is a need to sharpen our understanding of the intertwined analysis and guidance processes during mixed-initiative interactions in human-centered machine learning. In this paper, we introduce three novel components of co-adaptation in visual analysis. Figure 1 outlines these components as basic elements of co-adaptation at different levels of granularity: *learning phases*, *adaption processes*, and *interaction dynamics*.

At the coarsest granularity, we propose to model co-adaptation in distinct learning phases, each tailored towards a specific, measurable, and testable goal. Considering such goals and the associated learning phases facilitates reasoning about how users interact with the system and how adaptation can be made explicit. Furthermore, these goals provide structure for clearly specified and repeatable study designs. Inspired by the well-established Bloom's taxonomy [5] introduced in Section 3, we propose modeling the system adaptation over time based on three objectives: initialize, refine, and automate. These are general objectives that adaptive systems often strive to achieve that can be tailored to precise, testable goals for a given application. Furthermore,

*Corresponding author
e-mail: `fabian.sperrle@uni-konstanz.de` (Fabian Sperrle)

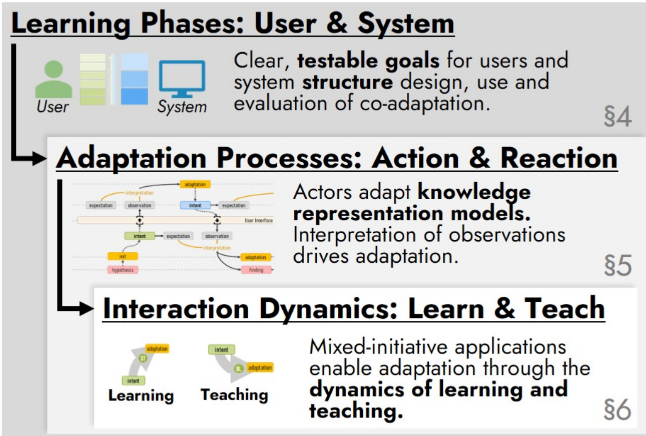

Fig. 1. Paper overview. Holistically considering learning phases (Section 4), adaptation processes (Section 5) and interaction dynamics (Section 6) structures the ongoing discussion about adaptive systems in visual analytics and reveals opportunities for future research (Section 8).

reaching a defined adaptation goal influences which interaction affordances are available and if and how often a system takes the initiative. An overview of the different learning phases and their human counterparts based on Bloom's taxonomy is provided in Section 4.

Adaptation processes are the second novel perspective on co-adaptation. Throughout learning phases, users and systems converge towards a common understanding of a shared task. To that end, they adapt knowledge representation models based on observations in mixed-initiative interaction. In Section 5, we characterize adaptation processes as action-reaction sequences that both actors observe. Actors interpret the observations with respect to their expectations and adapt when necessary.

The most fine-grained perspective of our model of co-adaptation is on interaction dynamics. Interaction dynamics structure user and system interactions into *learning* or *teaching* interactions. We elaborate in Section 6 how interaction dynamics provide a novel manner of reasoning about co-adaptation in visual analysis and are derived from the process-oriented view on adaptation. We postulate that both actors, i.e., the system and users, maintain implicit expectations towards their opposing actor's reactive behavior. Managing such expectations and communicating the capabilities of each actor influences the success of co-adaptive processes. We discuss such expectations and their impact on the co-adaptation model.

To conclude the paper, we provide various perspectives on the diverse design opportunities that the multigranular co-adaptation model offers across learning phases, adaptation processes, and interaction dynamics. As this conceptual model describes how co-adaptive processes can be designed and implemented, it opens up a space to reason about existing works. Hence, Section 7 highlights how we can structure interactive visual analytics techniques into the dimensions provided by our model. Based on this review of implementation examples, we derive research gaps and opportunities, as detailed in Section 8.

Overall, our main contributions are: (1) a multigranular model of co-adaptation spanning learning phases, adaptation processes, and interaction dynamics that places particular focus on making co-adaptation testable. (2) a discussion of actor expectations towards adaptation; and (3) an overview of research opportunities in co-adaptive analytics.

This paper extends our previous workshop submission on learning and teaching in co-adaptive guidance for mixed-initiative visual analytics [6]. We extend the previous work by four main aspects: First, we provide a background in educational theories and, in particular, Bloom's taxonomy. Second, we derive a set of goals for adaptive systems from Bloom's taxonomy. Third, we postulate a set of expectations that actors have towards the reactions of their opposing actors. Finally, we conclude by structuring the main research challenges for co-adaptation.

2. Related Work

Already in 1999, Horvitz described design principles for mixed-initiative systems, including "providing mechanisms for efficient agent-user collaboration to refine results" [1]. Further, Oppermann et al. [7] investigated adaptive systems that can respond to user input. In particular, early approaches describe the generation of "*knowledge bases*" for controlling adaptive dialog-based systems [8] and state that systems should model the user, the task, the domain, and themselves [9].

Mixed-Initiative Analytics. Mixed-initiative interaction has been described as "a flexible interaction strategy in which each agent (human or computer) contributes what it is best suited at the most appropriate time" [10]. However, in visual analytics, the concept of timing has often been neglected. Instead, mixed-initiative visual analytics systems were characterized by the inclusion of a recommender engine that suggested alternative visualizations (e.g., [11, 12]) or modeling alternatives (e.g., [13, 14]). Endert et al. have then argued "for a shift from a 'human in the loop' philosophy for visual analytics to a 'human is the loop' viewpoint, where the focus is on recognizing analysts' work processes, and seamlessly fitting analytics into that existing interactive process" [15], while Crouser et al. explored "how to balance the contributions of humans and machines in computational systems" [16]. Cook et al., instead, explored how to integrate task-driven recommendations into visual analytics based on initial, user-provided seeds representing entities of interest [17]. These systems are representative of the main challenges in mixed-initiative analytics: systems have to identify the right context and provide the right content (at the right time).

Guidance. Much more recently, user support under the name of *guidance* has become a topic of interest in the visual analytics community as a "promising attempt to enable a better collaboration of the human and the computer" [18]. Since then, three publications have shaped the definitions of guidance, in particular. Ceneda et al. [3] first characterized guidance in visual analytics in terms of an existing knowledge gap, available inputs and outputs, and the degree of guidance. They build on van Wijk's visualization model [19] to show where different types and degrees of guidance affect the model. Collins et al. criticize the model as "too abstract to use practically" [4]. They, instead, propose a model based on Andrienko et al.'s framework that characterizes visual analytics in terms of a model building process [20]. Building on this model specific to visual analytics enables the extraction of more concrete situations in which users might require help with typical tasks. Collins et al. state that the knowledge of an "intelligent guide" can be categorized as *prior knowledge*, *session-specific knowledge*, and *situation knowledge* [4]. However, they do not reason about how this knowledge updates over time and who has agency over the changes in effect. Federico et al. present a theoretical framework that incorporates "the function and role of tacit and explicit knowledge in the analytical reasoning process" [21]. Our work focuses on different learning goals that users and systems should reach over time and goes beyond knowledge collection towards application and synthesis.

Most recently, in their state of the art report, Ceneda et al. explicitly state that guidance is a mixed-initiative process and characterize existing approaches along the dimensions of user- and system-guidance [18]. Here, it is interesting to consider who initiated the guidance and who is adapting as a result. In human-machine collaboration, such adaptation processes have been studied [22, 23] and modeled game-theoretically [24].

In this paper, we focus on co-adaptation in mixed-initiative systems and provide an alternative view on guidance by considering learning and teaching processes. These processes are linked to the provision of explanations and should follow principles from pedagogy, such as clarity, elicitation of learners' responses, and relevance to the learner [25]. The relation of guidance to pedagogy will be introduced in more detail in Section 3.

Computational Steering and Model Steering. Computational Steering has been defined as "researchers [changing] parameters of their simulation on the fly and immediately [receiving] feedback on the effect" [26]. In visual analytics, this has often been named *model steering* and forms one of the cornerstones of the field. For example, model steering approaches exist for topic model optimization [27] and data exploration [28]. *Semantic interaction* is a special form of model steering in which the analytical reasoning is inferred from user interactions [29]. Semantic interactions constitute prime examples of the user providing guidance to the system, making them conceptually related to system-provided guidance. Our model of co-adaptation presented in Figure 4 models both cases as interactions with the intent to *teach*.

3. Foundations: Educational Theories and Learning Goals

Co-adaptation plays an important role in human-centered machine learning. In recent years, research has typically focused on integrating human knowledge into interactive machine learning processes in the first place. Now, human-centered machine learning places a particular focus on the work that humans perform [30], expecting systems to adapt over time to better support users in their tasks. However, as of now, there is no general evaluation framework that can help assess the success of human-centered machine learning in general or of co-adaptation in particular. As a result, it remains difficult to compare available systems and techniques and decide which approach should be employed in a given situation. To that end, we identify a need for characterizing what systems have learned and are still expected to learn in co-adaptive workflows.

Learning Theories. For human learners, various models of learning processes exist in education theory. In the context of human-centered machine learning, collaborative learning [31] and co-operative learning [32, 33] are particularly applicable. Broadly speaking, collaborative learning describes scenarios in which multiple actors work together to generate knowledge on a common subject [31, 34]. As a subset of collaborative learning, co-operative learning aims to address issues like team members that do not contribute to the final result by relying on five key properties: *positive interdependence, individual accountability, face-to-face interaction, interpersonal skill development*, and *assessment of team-functioning* [34]. While interpersonal skill development and assessment of team functioning are research questions closely related to psychology, positive interdependence and individual accountability are directly applicable in visual analytics, where users rely on the systems to be successful but are accountable for the overall result. In human-centered machine learning, this particularly means that humans must understand what systems are doing and how they are reaching their results. Current work in system intelligibility is a step in this direction and highlights that the intelligibility of "*datasets, training algorithms or performance metrics*" [35] could be more critical than traditional model intelligibility.

Blooms Taxonomy: Assessing Learning Goals. While collaborative and cooperative learning describe how teams should interact, they are not concerned with evaluating that learning goals have been reached. Instead, they often rely on Bloom's taxonomy [5] to verify that the learner understood and internalized the new knowledge. Such verification is also necessary in co-adaptive analysis, where it is difficult for users to know if, and what, systems have learned. As Schunk states, teachers "may believe that students have learned, but the only way to know is to assess learning's products and outcomes" [32].

To that end, Bloom's taxonomy defines six learning objectives that aim to measure the degree of understanding that a learner has achieved. The taxonomy is intended as a classification of behavior exhibited by learners and testable by teachers. During visual analytics, such testing can take the form of model probing or what-if analysis [36, 37] and allow users to verify their model.

Each objective in the taxonomy describes the expected outcomes from recalling facts and definitions (*remember* objective) to highly complex tasks such as analyzing information and assessing its value (*evaluation* objective). Clearly defined, distinct objectives for systems allow for a nuanced evaluation and present a first step towards a structured evaluation framework for co-adaptive systems. While the taxonomy originally describes hierarchical levels of understanding, Burns et al. [38] remark that this assumes a linear learning process that is not necessarily followed in practice.

In the visualization community, Bloom's taxonomy has been used directly or as inspiration in a diverse set of projects: to support designers in creating effective communicative visualization [39], to design study tasks [40], in a taxonomy for user engagement in information visualization [41], and, most frequently, for teaching (e.g., [42, 43]). Most related to our work, in their study to evaluate data visualizations, Burns et al. [38] implement questions targeted to each level of the taxonomy to measure the quality of their visualizations and complement conventional methods such as speed or accuracy testing. In their conceptual model of explanation processes in explainable artificial intelligence, El-Assady et al. [44] define verification blocks following each explanation block to ensure user understanding. Similar verification blocks could be added to co-adaptive analysis systems to ensure that adaptation goals are being met. In the following section, we elaborate on how Bloom's taxonomy integrates into visual analysis processes and how it can be adapted to verify system adaptation.

Explanation Process. In the field of pedagogy, learning and teaching processes are closely linked to providing explanations [25]. Odora concludes that effective explanations need to be context-aware and demand that teachers possess not only knowledge about the subject matter but also perceptive communication capabilities. Co-adaptive processes that aim to promote teaching and learning, therefore, need to follow the strategies and principles of adequate explanations and pedagogy. These include, amongst other things: clarity, eliciting a response from learners, and relevance to the learner [25]. Recently, research on the topic of explanations has been driven by the need for explainable artificial intelligence (XAI), as an increasing number of systems include fully automated decision making. Explanations are used to communicate information that the learner can use to improve their mental model and create expectations and predictions. This means that explanations are elicited from learners when they register situations that deviate from their expectations [45]. Miller [45] summarizes that although it is difficult to define what constitutes an explanation, explanations always refer to causality analysis. El-Assady et al. [44] present a conceptual model for the explanation process and explanation strategies. The explanation process can be defined as an iterative sequence of explanation-verification blocks. In each explanation block, three possible explanation strategies can be applied: *inductive reasoning*, explanation of detailed observations to be generalized to a bigger set of observations (bottom-up approach); *deductive reasoning*, explanation starts with a general overview and proceeds to more details (top-down approach); and *contrastive explanations*, comparison of two phenomena with sometimes implied contrast elements ("why not X?") [46]. All these characteristics of the explanation process show that a high degree of variability exists for the design of adequate explanations during co-adaptive analysis. Design decisions need to be based on different factors such as the target user group and their reasoning process [44]. Additionally, there are cases in which errors and bias need to be taken into account for explanations to help mitigate such errors [47].

4. Phases of Co-Adaptation: Users and Systems

Adaptation in visual data analytics is a continuous process, where an adaptation of the user can cause an adaptation of the system and vice versa. The overarching goal of this process is to reach a high degree of machine automation that remains intelligible and controllable through the user [48].

There are many different models of how to involve humans in a mixed-initiative visual analytics process: Sacha et al. introduced the knowledge generation loop [49], Liu et al. presented the problem solving loop [50], and Karer et al. provide a formal model of interpretation and reasoning in visual analytics [51], to name just a few recent examples.

We, instead, focus on how successful adaptation can be measured. As outlined in the previous section, several theories of collaborative and cooperative learning between (human) learners exist. Typically, they rely on *Bloom's taxonomy* [5] to define learning goals that should be reached over time. Such a classification is currently missing for adaptive systems. Explicitly considering learning goals for systems early on in the design process supports system developers in selecting effective interaction paradigms and aiming for system intelligibility. Furthermore, it represents a first step towards a unified evaluation methodology for adaptive systems as it provides comparable learning objectives.

To distinguish different interaction patterns employed to reach different adaptation goals, we map each goal to a *phase* that targets reaching the respective goal. While the adaptation goals for systems that we derive below are parallel to the goals for users and increase in complexity, the phases for system and user are independent. Hence, it is not a requirement that user and system are in "parallel" phases at all times.

In Subsection 4.1, we first briefly introduce Bloom's taxonomy that specifies learning goals (for human learners) across six levels of increasing complexity. We then map those existing learning objectives for users to a set of objectives for mixed-initiative systems in Subsection 4.2. These objectives reflect our experience from the design and evaluation of systems and are intended to spark a discussion on more clearly defined expectations towards adaptive systems. Finally, Subsection 4.3 outlines how the three proposed goals relate to existing models of human-machine interaction in visual analytics.

4.1. User Adaptation Phases

Learning in humans can be classified according to Bloom's taxonomy [5] across six sequential objectives. Burns et al. [38] provide a translation of six objectives (from an updated version of Bloom's taxonomy [52]) for information visualization that we summarize below. We build on their definitions and further adapt each definition towards visual data analysis.

Remember

Remember. The *remember* goal requires learners to recognize or recall previously learned "ideas, material, or phenomena" [5]. For visualization, appropriate tasks include locating and reporting specific pieces of information [38]. This definition is also applicable to visual analysis, although extended to include information about available system functionality.

Understand

Understand. The understanding goal states that learners should be able to interpret and extrapolate information, leading to "an understanding of the literal message contained in a communication" [5]. Usually, this is achieved by understanding or creating abstractions. Typical tasks in information visualization include the generation of data summaries or the generation of key takeaways [38]. In visual analysis, this extends to being able to understand different system functionality and data transformations and gauge their effect on the available data.

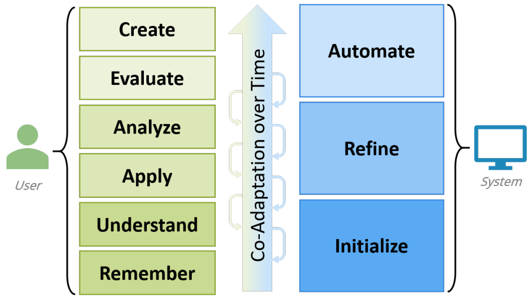

Fig. 2. Mapping the six phases of Bloom's taxonomy of learning objectives to three learning phases for co-adaptive systems. *Initialization* is about understanding the user's tasks and goals. *Refinement* analyzes user behavior and their reactions to fine-tune the model. *Automation* applies the learned information and allows systems to take the initiative in the analysis.

Apply. While *understanding* only requires learners to comprehend a certain abstraction and use it when prompted, *application* requires that the abstraction has been understood so well that it can readily be applied [5]. In visualization, appropriate tasks include the transfer of a visualization to a different visual representation [38]. Typical tasks in visual analysis include, e.g., the selection of a different machine learning technique or switching to alternative item-selection strategies in active learning.

Analyze. Learners that reached the *analyze* goal are expected to recognize unstated assumptions, identify and verify hypotheses, and find patterns [5]. Burns et al. translate this to the recognition of trends in the data or the identification of evidence to support specific conclusions [38]. For visual analysis, relevant tasks include the quantitative and qualitative evaluation of machine learning results and the identification of bias, as well as the creation and verification of hypotheses.

Evaluate. Learners are expected to make "quantitative and qualitative judgments about the extent to which material and methods satisfy criteria" [5] to complete the *evaluate* goal. For visualization, Burns et al. name the evaluation of a visualization with respect to given criteria or the provision of justification for given results as potential tasks [38]. In co-adaptive analysis, users should be able to judge both the quality of the obtained result and the quality of the system they used.

Create. Reaching the *create* goal of Bloom's taxonomy, learners should be able to combine their ideas and knowledge to form something "not clearly there before" [5]. Tasks in visualization include the prediction of further values in a sequence or the identification of different views that reveal new information [38]. Learners at this level are able to perform human-in-the-loop analysis, where the human drives the analysis by creating hypotheses and verifying or rejecting them to generate new information and knowledge.

4.2. System Adaptation Phases

A classification of learning goals such as Bloom's taxonomy is missing for adaptive, mixed-initiative systems. Building on the user goals for information visualization provided by Burns et al. [38], we propose a mapping to goals for adaptive systems. These goals for systems are derived from our previous work and reflect the current state of mixed-initiative systems. Figure 2 shows the mapping and highlights that both user and system typically progress towards more challenging learning goals over time. However, learning is a non-linear process. Consequently, goals could be reached out of order or can become outdated when information concerning previous goals changes. Nonetheless, these goals can structure how we reason about mixed-initiative systems. For example, in the design process, they require system designers to consider if and how system behavior can change once a goal has been reached. Similar to a teacher who does not rely on the same teaching paradigm across all learning goals, different interaction paradigms might be more appropriate at different times. Furthermore, clearly defined goals for systems provide an opportunity for more comparable system evaluation to verify that goals are met and understand the users' perception of systems in those different stages.

Initialize. The initialization objective combines the two initial goals *remember* and *understand* from Bloom's taxonomy. To complete this objective, systems should **observe what data the user is working with** and **identify potential user tasks** that they might perform (remember). Additionally, they should **create initial models** of the observed information, demonstrating their ability to identify appropriate abstractions (understand).

In the context of mixed-initiative systems, systems in this phase should refrain from taking the initiative to avoid irritating users. Instead, they should focus on observing user interactions and building models that are able to capture intent. The time needed to complete this stage can be significantly shortened by pretraining appropriate models outside of the mixed-initiative interaction loop.

Refine. After the system has created initial knowledge representation models (see Subsection 5.1) in the initialize phase, the overarching goal in this phase is the *refinement* of those models to fit individual users and their tasks. More specifically, systems should **identify in which situations users perform which tasks**, how to **characterize their analysis behavior**, **capture which parts of the data are most interesting**, and where a **user might encounter knowledge gaps**. As such, this phase is closely related to goals of guidance in visual analytics [3]. Here, adaptation induced by the dynamics of learning and teaching is most prevalent.

Systems providing active guidance should begin to make initial suggestions in this phase and gather user feedback in order to improve their future suggestions. Most system-initiated actions in this phase should have a guiding nature, as systems are expected to still refine their knowledge representation models, leaving them unable to provide good automated analysis steps.

Automate. Once the system has reached a sufficiently accurate model in the refinement phase, it can begin to *automate* expected user actions. This increased initiation of actions frees the user to work on other tasks and can increase the efficiency of the analysis

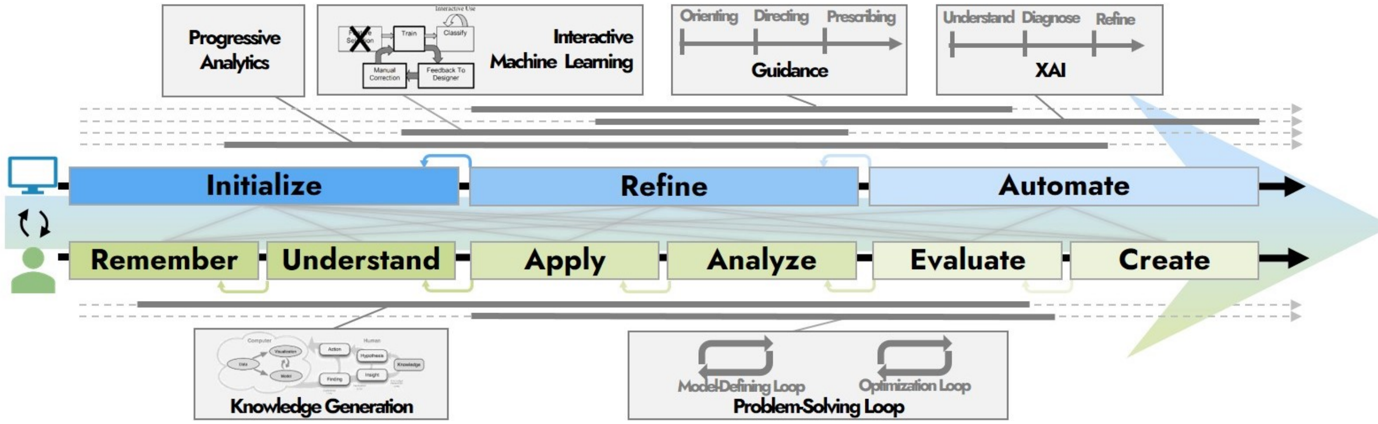

Fig. 3. The co-adaptive analysis process traverses the three phases *initialization*, *refinement*, and *explained automation*. The figure shows relations to related concepts and processes, connecting established models to our proposed co-adaptive process. The models we consider are Progressive Analytics [53], Interactive Machine Learning [54], Guidance [3], Explainable Artificial Intelligence (XAI) [55], Knowledge Generation [49], and Problem-Solving Loop [50].

Automate

process. System goals at this level include **advancing the analysis**, **automated evaluation of analysis results** (e.g., via learned relevance metrics), or the **suggestion of alternative analysis strategies**. When designing systems that aim to reach this level of adaptation, designers should include overview and verification phases to ensure that users can still review system actions and potentially intervene, leaving them with control over the process and a sense of agency.

4.3. Relation to Existing Models

There are various different models that describe the visual analytics process or relevant related concepts and techniques. Some of them, like *guidance*, are actively used in co-adaptive analysis systems. Others, such as *knowledge generation* or *interactive machine learning*, focus on adaptation in the human or the system, respectively. However, they can all be used to describe analysis processes over time. Figure 3 illustrates that our model of co-adaptive analysis, in general, and the three levels of learning goals for systems, in particular, can represent those related models on an abstract level. The arrow, again, displays that co-adaptive analysis processes tend towards more system-initiated actions over time. However, similar to learning in humans (Figure 2), this is not a linear process, and systems must be able to return to previous objectives or complete objectives out of order. The figure shows that many existing theories and models cover similar adaptation processes under different names or go through similar phases of adaptation. A unified structure of learning goals can enable the effective comparative evaluation of systems employing those different paradigms. We believe that our initial set of learning goals is intuitive and easily applicable to a multitude of systems.

Knowledge Generation. The knowledge generation loop [49] as a generic model of insight generation can span the entire range from human initiative to system initiative (although typical implementations primarily use human initiative). The exploration loop addresses the analyst's goals of getting to know (remember) and understanding the data. The system, on the other hand, can initialize during this phase. The verification loop is used to "confirm hypotheses or form new ones" [49], matching the learning goals *apply* and *analyze*. Systems can use this phase to refine their models. In the knowledge generation loop, users evaluate whether to trust their insights and, if yes, generate knowledge.

Problem Solving Loop. The problem solving loop [56] contains two sub-loops. First, the model-defining loop builds on mathematical representations of the problem created by experts. Consequently, the initialization does "not need to be supported by an interactive optimisation tool" [50]. A system can, however, support users by aiming to *refine* the model by adding, removing, or changing constraints and objectives of the optimization model [50]. In the optimization loop, users then *analyze* different solutions before *synthesizing* and *evaluating* them. In this phase, systems can support users by exploring alternative optimizations and presenting ranked lists. However, the final decision on solution quality is a user task. Thus, systems do not have to fulfill all goals of the *automate* objective.

Interactive Machine Learning. Endert et al. [54] surveyed the state of the art in integrating machine learning into visual analytics. They categorize the field along the two primary tasks of modifying parameters and defining analytical expectations. Both tasks primarily fall into the *refine* and *apply/analyze/evaluate* levels for systems and users, respectively. This assumes that users already have a basic understanding of their data and that the optimized machine learning model has been initialized outside of the mixed-initiative process.

Guidance. "Guidance is a computer-assisted process that aims to actively resolve a knowledge gap encountered by users during an interactive visual analytics session" [3]. Guidance systems can profit from having access to a user's interaction history and previous analysis states. Consequently, guidance systems should observe users from the beginning of the analysis to *initialize* their models. Once they have gathered a sufficient understanding of the data and the user's task, they can begin to make initial suggestions and observe the reaction to *refine* their models.

Interactive Model Analysis for XAI. For explainable artificial intelligence, Liu et al. [55] define a three-phase model called *interactive model analysis* to understand, diagnose and refine machine learning models with the help of bespoke visual analytics techniques. Users should be enabled to *understand* "why machine learning models behave the way they do" [55]. Further, they should analyze potential training failures before accepting *automated* guidance from the system to refine their models.

Progressive Analytics. Progressive analytics is not a mixed-initiative process by definition. However, Fekete et al. [53] describe a three-phase model of uncertainty in progression. In the first phase, called *estimating suitability*, the model is just beginning to process data. In the second phase, it provides an *early response* that becomes more and more certain over time. In the third phase, the model converges, and its uncertainty has stabilized, allowing the user to advance the analysis.

5. Processes of Co-Adaptation: Actions and Reactions

This section introduces the process of co-adaptation over time that is shown in Figure 4. In this process, both user and system adapt knowledge representation models based on their mixed-initiative interaction. The co-adaptation process operationalizes how the learning goals introduced in Section 4 can be reached.

In Subsection 5.1 we first introduce several knowledge representation models that adapt over time. We then define key concepts of our process of co-adaptation and highlight connections to previous work in Section 5.2, before introducing the iterative progression of co-adaptation in Subsection 5.3.

5.1. Knowledge Representation Models

Before considering the *adaptation* of a system and user, we establish the types of knowledge and information both actors have. Knowledge and information relevant to the analysis process are stored in knowledge representation models.

Krogsæter and Thomas state that knowledge-based systems require models of the user, the task, the domain, and themselves (system model) [9]. According to their definition, the system model should contain knowledge that the system has about its functionality and limitations. As this information is unlikely to change during the mixed-initiative analysis process, we do not consider it. Instead, we define system models as representations of the user's knowledge about the system.

Data Model. The data model contains information such as data distributions, descriptive statistics, identified outliers, and relations and similarities between data points. Typically, systems are expected to have a complete data model due to their increased computational abilities.

User Model. The system stores a specific user model for each user. This model contains all knowledge that the system has explicitly or implicitly gathered about the user. The user model aims to capture, among others, the users' knowledge, their level of expertise, potential biases, personal preferences, and personality traits. Beyond knowledge, user models should also consider the user's cognitive abilities such as perceptual speed, visual working memory, and verbal working memory, as personalization can counteract these inter-user performance differences [57].

System Model. The system model is the mental model of the system that users create during the analysis. It includes knowledge about the implemented algorithms with their strengths and weaknesses, available visualizations, and guidance operations that the system offers. The system model is created over time through interaction with the system and influenced by previous knowledge of similar systems. The system model, therefore, fundamentally influences the expectations the user has about each task outcome.

Task Model. The task model contains all necessary knowledge to solve the tasks along the analysis process, including the order of task execution, the (hypothesized) solutions, relations and similarities between tasks, and the analysis context.

5.2. Components of Co-Adaptation Processes

Before presenting the co-adaptation process in Subsection 5.3, we define all terms used and provide relations to relevant previous work where applicable.

Action. We follow the definition by Gotz et al. [58] and define actions as aggregations of semantic sequences of individual events. As those actions that are relevant to future adaptations in the process have an associated intent (see next paragraph), we omit individual actions from the visual representation of the process in Figure 4.

Intent. In the context of mixed-initiative analytics, there might be different types of actions, e.g., some that are specific to advancing the analysis and others that are meant to refine future adaptation. The goal of each action is its *intent*. As our model is applicable to co-adaptive visual analytics in general, we do not systematically differentiate between guidance intent and analytical intent.

Wenskovitch et al. model the complex relationship between interactions and intent [59]. They identify four different types of relationships: *One interaction implies one intent* (e.g., direct manipulation of a slider), *many interactions imply one intent* (e.g., different ways to switch fonts in Microsoft Word), *one interaction implies many intents* (e.g., moving data points in a projection view) and *many interactions imply many intents*. For a more thorough description, see Wenskovitch et al. [59].

Expectation. Expectations are closely related to the intent introduced above. They capture the predicted reaction to a given action and what adaptation a given action should introduce.

Adaptation. We define adaptation as the sum of all changes induced in the models (data, task, user/system) described in more detail in Subsection 5.1. The adaptation is the result of an interpretation of the observed actions and the derived intent. Similar to the four interaction-intent relationships introduced above, there exist four analogous intent-adaptation relationships.

Observation. Observations of the system-side are any recognized inputs that the system processes. An observation on the user-side is a (typically) visual change to the user interface (e.g., an updated model visualization, a log message, or any other form of perceivable change to the system state).

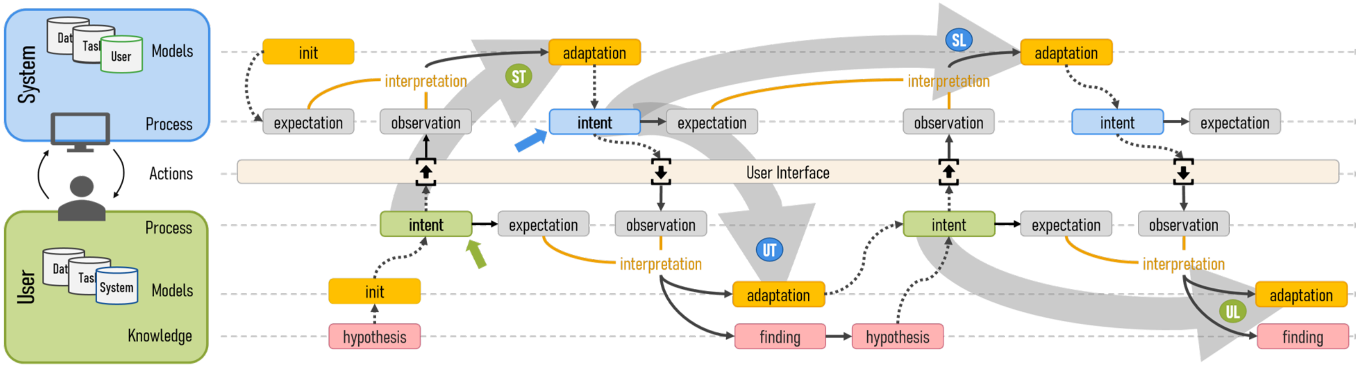

Fig. 4. Actions and reactions between user and system form the foundation of the co-adaptive analysis process. Reactions are observed and compared to an expectation, leading to the adaptation of the data, task or user/system models, and the derivation of new intents. Here, the user initiates the process (green arrow), and both the user and the system adapt. The system can also initiate the process, which would then start at the blue arrow. The grey arrows indicate the learning and teaching dynamics: system teaching (ST), user teaching (UT), system learning (SL), and user learning (UL). Figure by Sperrle et al. [6].

Interpretation. Both systems and users interpret the observed inputs in combination with intents and expectations. For the system, this typically means some form of machine learning and will be discussed in Section 7. Users often have to interpret changes to a visual representation instead as they cannot typically probe the system's models directly. In order for this interpretation to be possible, system adaptation must trigger an appropriate visual change that matches the user's expectation.

Kindlmann and Scheidegger provide a theory of algebraic visualization design [60]. According to their framework, the amount of change to a visualization must be proportional to a change in the underlying data. In co-adaptive analytics, this principle should not only apply to visual representations of the system's state but also inform the size of behavioral changes, with consequences that are predictable for a user. "Small" interactions (e.g., the labeling of few data points) should only lead to small, incremental changes in behavior. In Subsection 6.3, we outline user expectations towards system changes in more detail.

Hypothesis. A hypothesis "formulates an assumption about the problem domain that is subject to analysis" [49]. It can form the starting point for a co-adaptive analysis cycle and directly influences the performed actions. However, not every visual analytics task requires a hypothesis, and new intents could, e.g., also be based on interpretations of previous observations.

Finding. We again rely on the definition by Sacha et al.: "A finding is an interesting observation made by an analyst using the visual analytics system" [49].

5.3. Interactive Progression of Adaptation

The goal of visual analytics is to incorporate human intuition in the analysis process to generate hypotheses and extract knowledge [49]. As a result, the user's questions might change during an analysis session. For example, users might become aware of unexplored regions of the data or additional system functionality that could be beneficial to solving the current task. Systems may capture the task users are trying to solve more accurately. Both user and system need to adapt over time and take the progression analysis state into account to accommodate this change.

Interaction and Adaptation. We provide a detailed model of interaction and adaptation in co-adaptive analysis processes in Figure 4. The model shows an interactive analysis process, where the x-axis represents time. It is centered around action-reaction pairs ↥→↧ that are exchanged between the user and the system. Building on the analytic activity model by Gotz et al. [58], we define *actions* as aggregations of individual events. Actions are, in turn, aggregated into one or multiple higher-level user or system *intents* (see Section 5.2). Each intent is associated with a corresponding expectation that captures the assumed impact of the performed actions. We are particularly interested in those expectations that concern changes to the (mental) models of the recipient of performed action(s). Users and systems *interpret* their *observations* and expectations with respect to the available models (data, task, and user/system) and, in the case of the user, knowledge. The result of this interpretation may lead to an *adaptation* of the recipient, as well as the generation of new *findings*.

Opportunities for Adaptation. Figure 4 shows an interaction in which all opportunities for interpretation and adaptation have been realized. In practice, many actions will not be interpreted, e.g., because most current systems lack support for intent identification, and users might choose to focus on their task at hand rather than analyzing every system action. Figure 4 also shows grey arrows that indicate four dynamics that drive adaptation in co-adaptive interaction: *system teaching*, *system learning*, *user teaching*, and *user learning*.

6. Dynamics of Co-Adaptation: Learning and Teaching

The success of co-adaptation for solving high-level analysis tasks depends on interaction dynamics and expectations that actors have towards them. The co-adaptive analysis process model in Figure 4 reveals the two central interaction dynamics of learning and teaching. In this context, we define the actors' intent to learn as the aim to adapt themselves, with the help of knowledge provided by another actor. Conversely, we define the intent to teach as the aim to induce adaptation in the other actor. As both users and systems can initiate both learning and teaching, there are four different dynamics that provide a process-oriented view on interaction dynamics in co-adaptive analysis: *user teaching*, *system teaching*, *system learning*, and *user learning*. Our focus is on the adaptation that is caused in a given actor. Hence, we reference the adapting actor in the names of the four dynamics. *System teaching*, e.g., describes a dynamic in which the user provides knowledge to the system, causing it to adapt. It is important to note that neither actor adapts in isolation. Instead, the feedback from the other actor is fundamental in providing knowledge and resolving the encountered knowledge gap. Consequently, *system learning* is different from general *machine learning*.

The four dynamics are responsible for the adaptation in the initialization and refinement phases. During initialization, systems typically have to rely on *system teaching* provided by the user. During refinement, systems can begin to require learning guidance as well. During the analysis process, these dynamics often do not appear in isolation but can be interleaved, as Figure 4 illustrates. Ultimately, systems should aim to enable multiple, if not all, dynamics if they are to be mixed-initiative systems. In the following section, we briefly introduce each of the dynamics in more detail before describing expectations towards interactions that must be considered to successfully enable them. The figures presented with each dynamic are excerpts from Figure 4 and locate the dynamic in the interaction process. For real-world examples that represent these principles, see Subsection 7.2.

6.1. Teaching Intent

The *teaching dynamic* is initiated by an actor that aims to adapt the models of the other actor. Goals for teaching include providing help in a given situation to facilitate the analysis, informing about alternative analysis options, suggesting potential corrections, explaining the current model, or providing a tour as guided exploration. Typically, system-provided teaching targets the data and task models of users. In contrast, users typically teach systems about the task and their subjective preferences.

User Teaching. *User Teaching* is the most commonly implemented in the form of guidance in modern systems, where systems aim to teach users. It directly translates to the original goal of guidance, which is resolving encountered knowledge gaps. To that end, systems, e.g., highlight data points to consider [61], or present recommendations and alternative analysis pathways [62].

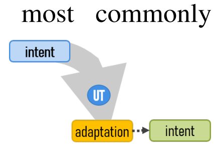

System Teaching. In *system teaching*, the user aims to teach the system their understanding of the task or data. As such, it is closely related to the concept of machine teaching [63]. However, while machine teaching is typically concerned with providing systems with "labels, features [or] structure" [63], system teaching also allows systems to update their user model with, e.g., observed preferences and biases.

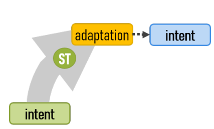

6.2. Learning Intent

Actors request *learning guidance* with the intent of verifying or adapting their own models. Beyond asking for help with the analysis, the goals of learning guidance include probing the other actor's models, verifying hypotheses, and understanding the current situation.

User Learning. Users initiate *user learning* with the goal of learning about the data, the system, or its understanding of tasks. This operation can be considered a *probe*, providing users with additional knowledge and not necessarily advancing the analysis.

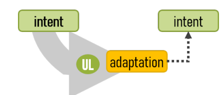

System Learning. *System learning* describes guidance in which the system requests user feedback with the aim of improving its user, task, or data model. While this operation may or may not have an immediate benefit to the analysis process, the gathered information can be used to improve further guidance as it helps systems to understand users and tasks better.

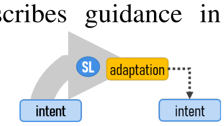

6.3. Expectations for Interactions and Adaptation

Both learning and teaching interactions can be used to adapt the knowledge representation models of either actor and only differ in the initiator of the adaptation. Selecting the correct dynamic at the correct moment has a large impact on the success of mixed-initiative systems as it impacts the perceived agency [64]. However, other factors beyond agency and locus of control affect not only result quality, but also user satisfaction in co-adaptation workflows. In particular, there are several explicit or implicit expectations that actors have towards interactions and the adaptation that they induce. These expectations exist for both the system and the user: on the system side, expectations form the basis for assumptions in the used machine learning models. For example, designers of a learning-to-rank model might assume that users are bad (or good) at providing ratings that form consistent transitive relationships. On the user side, their expectations will drive their interaction patterns and lead to frustration if not respected.

Guidotti et al. [65] provide a survey of methods for explaining black-box models and derive several desiderata for interpretable models from the analyzed state-of-the-art. While they interpret the desiderata like accuracy and consistency with a specific focus on machine learning models, we argue that some can be reframed to describe expectations that actors have towards interactions in co-adaptation. Below we present an initial set of expectations towards adaptation. We ground those expectations on the desiderata for interpretable models as compiled by Guidotti et al. [65] and frame them from a perspective of mixed-initiative interaction. Future research in co-adaptive analytics should investigate which other expectations exist, how they can be identified from user behavior, and what the implications of violating those expectations are.

Monotonicity. Previous work has found that users are more likely to trust and accept classification models when they are built respecting the monotonicity constraints expected by the users and the domain [66]. This suggests that users might also expect systems to infer monotonic constraints from their interactions rapidly. This becomes a particular challenge if user interactions do not appear monotonic to the system, e.g., because it learns on a representation that does not match the user's mental model.

Accuracy. In XAI, accuracy is defined as a measure of the extent to which the model predicts unseen instances [65]. In co-adaptive analytics, accuracy describes how good the model is at identifying characteristic features from vague, semantic interaction. The more accurately the system can infer attributes, the more accurately its adaptation will match user intent.

Causality. A controlled change in the input data should affect the model behavior [65]. In co-adaptive analytics, each adaptation of the user or the system should be related to a change originated by the other actor. Causality is a fundamental property that enables the *interpretation* of observed changes.

Correctness. Typically, both the system and a user will assume the actions of the other actor to be correct. However, both actors can make errors of varying severity. Users could accidentally misclick or provide wrong information to the system due to unawareness. Systems could provide wrong information as a result of implementation bugs or due to biased training data. As errors are often unavoidable, this strongly relates to how a system recognizes accidental errors and possibly helps mitigate them.

Consistency. In co-adaptive analytics, the knowledge representation of both the system and the user can be seen as predictors of the other actor's action. Assuming the correctness of those models, an actor is consistent when it complies with the model.

Transitivity. Preference relations derived from user interactions are likely not transitive. Tversky argues that people represent objects as collections of features and potentially rely on different subsets of features when making pairwise comparisons [67]. As a result, he argues "that similarity, as one might expect, is not transitive" [67]. Consequently, systems should take care not to rely on transitive information derived from user inputs. Similarly, they should make learned transitive relations explicit. More generally, Tversky's findings suggest that system adaptations might not be easily transferable between users. However, this limitation is not specific to transitive learnings but applies to any domains and tasks in which there is no single correct answer and user preferences play a role in the final outcome.

Agility. Both actors, but especially systems, must be able to adapt to changing behavior quickly. Langley [68] finds that systems relying on user input, and thus the users' time, should rely on "induction methods that achieve high accuracy from small training sets over those with higher asymptotic accuracy but slower learning rates". They express the expectation that "an adaptive interface that learns rapidly should be more competitive than ones that learn slowly." [68]

Generality and Transferability. For machine learning models, Guidotti et al. state that "it is preferable to have portable models that do not require special training regimes or restrictions" [65]. The same is true for knowledge representation models. From our experience, novice users often expect systems to learn abstract and well-generalized knowledge.

7. Co-Adaptation in Existing Applications

Co-adaptive analysis systems learn to adapt the user, tasks, and data models as a result of user interactions with the system. From these interactions, different levels of input are available for model computation: low-level interactions, such as mouse movements, and high-level interactions, such as manual filtering actions. These types of interactions imply differences in data quantity, certainty, and continuity. For example, we can assume that mouse movements are recorded continuously throughout a session leading to many more data entries compared to punctual user clicks. Algorithms employed to classify, recognize and predict user interactions based on these different types of inputs, therefore, operate differently. Endert et al. [54] published a state of the art report on machine learning implementation in visual analytics describing important categories of algorithms used in the field. Co-adaptive systems predominantly need to perform classification, clustering and regression analysis on static and streaming data. Prominent algorithms used for adaptation in visual analytics systems are decision trees [69], Naive Bayes classification [70, 71, 72, 73], Hidden Markov models [74, 75, 76], nearest neighbor search [77, 78, 79], neural networks (e.g., self-organizing maps [80]), active learning [81], and reinforcement learning (e.g., Q-learning [82]). However, Endert et al. [54] note that new algorithms specifically tailored to incorporating user interaction into prediction are needed.

The previous sections have introduced both learning goals for adaptive systems and characterized adaptation in terms of learning and teaching. In this section, we present examples of applications of the proposed theoretical models in existing systems. As these systems do not typically mention these concepts explicitly, we rely on our interpretation of the system descriptions. Evaluating and reporting these human-centered factors of visual analytics represents an opportunity for the community working on co-adaptive systems.

7.1. Navigating the Phases of Co-Adaptation

We have presented a set of learning phases in co-adaptive systems. Here we introduce existing systems that follow our understanding of these learning phases and represent real-world examples for tangible implementation ideas.

Initialization. Micallef et al. [83] implement a method in which the system first iteratively asks for user feedback to learn task-based feature relevance until a certain quality is attained and subsequently generates predictions. This approach models the user's knowledge relevant for their specific task by adding upper confidence bound criterion computation [84] to the linear regression prediction model. They show that this initial phase of data gathering through user input is effective for increasing prediction accuracy in their user study.

Refinement. Many co-adaptive systems implement a refinement phase, with some directly entering this phase without going through co-adaptation for initializing and instead relying on assumptions, mathematical models, or external pretraining. Ottley et al. [74] implement an approach that directly permits refinement using Hidden Markov models that do not need prior data gathering and allow for a direct prediction of the user's next interaction. Reda et al. [85], however, observe that a limitation of Markov Chain-based approaches is the memoryless nature, with models losing the ability to predict high-level strategies. In fact, Monadjemi et al. [73] implement a system employing a Naive Bayes classifier that outperforms the Hidden Markov model approach. In VIANA [86], a system for argument annotation, Sperrle et al. eliminate the need for initialization by relying on the output of a domain-specific rule-based pre-annotation as initial suggestions. During the refinement phase, the user then promotes or discourages the system to suggest more similar or dissimilar annotations.

Automation. The system PlotThread [82] implements reinforcement learning to teach an AI agent how to draw storyline visualizations. After a refinement phase where users draw their own visualizations and refine them through interactions on the visualizations, the model learns how to generate and enhance visualizations based on the learned features. In their evaluation, the authors show how this method achieves better results at the cost of more iterations compared to a greedy algorithm that randomly selects interactions.

7.2. Employing Learning and Teaching Dynamics

In Section 6, we have explained how co-adaptive analysis includes the concepts of learning and teaching. We present example implementations of these dynamics in existing systems.

User Teaching. Shao et al. [87] support users during the exploration of large scatter plot matrices: based on eye-gaze data, the system shows plots that are visually dissimilar from those already explored. This guidance aims to teach users an unbiased data model that considers all data regions and maximizes the amount of information analyzed per time interval. A similar approach has been used by Silva et al. for gaze-based pattern recommendation [88]. LightGuider is a VA application for creating lighting designs [14]. Here, user teaching supports users in efficiently exploring the large model parameter space, enabling faster task completion by providing alternative model parametrizations while still supporting "manual intervention and artistic freedom" [14]. NEVA [89], a system for fraud detection in consumer networks, supports users during navigation of subgraphs to avoid non-plausible queries.

System Teaching. In current applications, system teaching is typically realized via explicit user inputs: users adapt target sliders [14] or create new entity relations [90]. Podium, a system for ranking multivariate data, includes guidance from the user [91]: users teach the system their understanding of the relations between data records by reordering them in a table. The system then infers a feature weighting model, capturing "*which attributes contribute to a user's subjective preference for data*" [91]. As the model is transparently made available to users, they can compare expectations and observations to make changes.

User Learning. Clustrophile 2, a system for interactive cluster analysis [69], suggests algorithms with various parameter settings. Users can ask for support from the system during feature selection or algorithm parametrization by toggling the *Help me decide* menu. The system provides, e.g., feature relevance scores or silhouette coefficients for selecting the number of clusters.

System Learning. Micallef et al. [83] developed an application that supports users during the generation of machine learning models with small data sets. The system employs a user model and asks users to refine features in a subset of the overall features by assigning user relevance for the overall prediction task. This step is initiated by the system to learn the user's domain knowledge, repeating the knowledge elicitation step as many times as necessary until the prediction model returns improved predictions. Further approaches include feedback-driven view exploration [77] and DataTone [92]. The system BEAMES [79] elicits feedback from users to update sampling weights for prediction models in the recommended models pool.

7.3. Assessing Co-Adaptation Expectations

We explore examples of systems that consider important expectations of actors towards adaptation in their approach. We show how these systems tackle the problem in their concrete use case. It is worth noting that none of these systems evaluate these expectations (see Section 8).

Causality. Actors learn if they can recognize a causal relationship between interactions and adaptations (see Section 3). In the system LightGuider [14], updating preference weights for light constraints triggers an update of the provenance tree, which shows how different simulated lighting scenarios reflect given preferences. Each update recomputes the "usefulness of each action for reaching a certain illumination constraint" using a Weighted-Sum Model [93]. Explicit user interface controls allow for a direct link between action and adaptation, but more implicit mechanisms can also be implemented for continuous causality inference. In such cases, providing explanations for the adaptation is required [94].

Consistency. Once the system learned user preferences, future actions (e.g., predictions or suggestions) should be consistent with the learned model. In BEAMES [79], previously saved model types are included in model recommendations for subsequent analysis stages by increasing the sampling probability for saved models.

Correctness. Actors rely on learning truthful information from each other, but oftentimes inputs can be erroneous or uncertain. To mitigate this, more advanced techniques include recognizing bias to show where misconceptions might lie using Hidden Markov Models [91] or Bayesian networks [73]. Healey and Dennis [72] model user interest using a boosted Bayesian network classifier [95]. They include uncertainty in the model adaptation from implicit user input by including an additional "uncertainty weight" to the boosting process.

8. Research Opportunities

We have proposed a process model of co-adaptation in visual analysis and mapped existing learning goals from research in pedagogy to learning goals for adaptive systems. To demonstrate our model's applicability, we have discussed how a multitude of visual analytics approaches can be described by our proposed process model. We see much potential for structured, qualitative evaluation to advance the field of co-adaptive analytics. This section highlights the most promising research opportunities based on our analysis of the current state-of-the-art.

8.1. Structuring the Design Space of Evaluations

A structured analysis of the design space for evaluations is necessary to mature the field of co-adaptive analytics. Such a design space would allow both the systematic evaluation of learning and teaching dynamics and a review of phases of adaptation. As a first step towards this goal, future work should survey existing work on co-adaptation and identify all related fields and communities, including machine learning, human-computer interaction, information visualization, and psychology. Especially the HCI community has a long history of modeling and evaluating adaptive systems. We expect that the different communities focus on a different aspect of co-adaptation. The machine learning community does not typically involve users in their evaluations, while the HCI community focuses on presentation and interaction design. Synthesizing the results from the different communities can reveal gaps in the evaluation of co-adaptation. Future research should then investigate which effects that exist in isolation can successfully be integrated into full-fledged analysis systems.

8.2. Verifying Expectations for Co-Adaptation

In Subsection 6.3, we have presented an initial set of properties that both systems and users expect in co-adaptive processes. They are derived from research in explainable artificial intelligence and interactive machine learning. To the best of our knowledge, these expectations have not yet been evaluated in the context of co-adaptation. We encourage researchers from visual analytics to collaborate with psychologists and HCI researchers to evaluate which expectations are most important to users and whether similar users have similar expectations. We expect that violating those expectations affects user satisfaction, trust in the model, and perceived transparency and interpretability [96].

However, the results of those investigations do not only impact user satisfaction. They also provide valuable input to system designers aiming to develop new co-adaptive applications. Making the right assumptions about user interaction patterns is crucial to avoid deriving wrong or misleading information during co-adaptation. We expect that several iterations of calibration might be necessary until a hypothesized expectation can be successfully included in a co-adaptive analysis system. Providing design studies of both failed and successful attempts will provide helpful orientation for system designers.

8.3. Understanding and Probing Co-Adaptation

The co-adaptation model presented in Subsection 5.1 surfaces how systems and users converge towards a common analysis process over time. In a first step, systems designed for co-adaptation should make adaptation observable to enable interpretation, e.g., through provenance visualizations. When a provenance visualization is not practical or available, systems should provide ways in which users can effectively probe the model to verify adaptation. One possibility to enable such probing interactions are model sandboxes that users can interact with in isolation [37]. Those sandboxes should ensure that users can freely explore what the model has learned without fear of breaking the model or biasing the model's future learning through their exploration.

8.4. Controlling and Steering Co-Adaptation

In addition to enabling the observation and probing of adaptation, system designers need to ensure that both users and systems can steer and control the adaptation. Intuitively, this means that users must keep control over the analysis process and should be given direct interaction possibilities to refine system behavior when automatic adaptation fails. However, future research should also investigate how the quality of user input can be assessed and whether "bad" user input can be rejected. Such an assessment could, e.g., be performed by comparing the input to other users' input or by computing quality metrics such as precision and accuracy.

Currently, integrating co-adaptation into visual analytics systems is a resource-intensive, bespoke process. In the future, system implementations could be significantly simplified through general libraries and frameworks that perform typical actions, like observing where users click or which views they utilize most often. In information visualization, frameworks like vega-lite [97] make the creation of visualizations significantly simpler. Similar frameworks for co-adaptation could provide a starting point for the rapid prototyping of co-adaptation.

9. Conclusion

In this paper, we have presented a multigranular model of co-adaptation in visual data analysis and guidance processes. To structure adaptation over time and promote the definition of testable adaptation goals, we proposed a three-level taxonomy of learning objectives for adaptive systems derived from Bloom's taxonomy of learning objectives from pedagogy. To clarify how actors adapt based on observed sequences of actions and reactions, we presented a process model for adaptation and characterized interaction dynamics in terms of learning and teaching. In addition to these interaction dynamics, we have also identified user and system expectations towards interactions and adaptation in co-adaptive systems.

We have shown the model's applicability in application examples from recent works in visual analytics. As described in our opportunities section, our future work will constitute implementing a complete system that incorporates co-adaptation as its leading design paradigm. We intend to conduct a systematic evaluation of the influence of the different dynamics levels, dynamics, and expectations on the co-adaptation process. Based on the findings we derive from such a system design, implementation, and evaluation, we aim to devise and collect successful strategies for co-adaptation in mixed-initiative systems.

Acknowledgements – This work has been partially funded by the DFG within grant number 455910360 (SPP-1999).

References

[1] Horvitz, E. Principles of mixed-initiative user interfaces. In: Proc. SIGCHI Conf. on Human Factors in Computing Systems. ACM; 1999, p. 159–166. doi:10.1145/302979.303030.

[2] Keim, D, Andrienko, G, Fekete, JD, Görg, C, Kohlhammer, J, Melançon, G. Visual analytics: Definition, process, and challenges. In: Information visualization. Springer; 2008, p. 154–175.

[3] Ceneda, D, Gschwandtner, T, May, T, Miksch, S, Schulz, HJ, Streit, M, et al. Characterizing Guidance in Visual Analytics. IEEE Trans on Visualization and Computer Graphics 2017;23(1):111–120. doi:10.1109/TVCG.2016.2598468.

[4] Collins, C, Andrienko, N, Schreck, T, Yang, J, Choo, J, Engelke, U, et al. Guidance in the human-machine analytics process. Visual Informatics 2018;2:166–180. doi:10.1016/j.visinf.2018.09.003.

[5] Bloom, BS, Engelhart, MD, Furst, EJ, Hill, WH, Krathwohl, DR, et al. Taxonomy of educational objectives: the classification of educational goals. Handbook I: Cognitive Domain. McKay Company, Inc.; 1956.

[6] Sperrle, F, Jeitler, A, Bernard, J, Keim, DA, El-Assady, M. Learning and Teaching in Co-Adaptive Guidance for Mixed-Initiative Visual Analytics. In: EuroVis Workshop on Visual Analytics (EuroVA). The Eurographics Association; 2020,doi:10.2312/eurova.20201088.

[7] Oppermann, R, editor. Adaptive user support: ergonomic design of manually and automatically adaptable software. L. Erlbaum Associates Inc.; 1994. ISBN 978-0-8058-1655-6.

[8] Tißen, A. Knowledge bases for user guidance in information seeking dialogues. In: Proc. Int. Conf. on Intelligent User Interfaces. ACM; 1993, p. 149–156. doi:10.1145/169891.169945.

[9] Krogsæter, M, Thomas, CG. Adaptivity: system-initiated individualization. In: Adaptive user support: ergonomic design of manually and automatically adaptable software. L. Erlbaum Associates Inc. ISBN 978-0-8058-1655-6; 1994, p. 67–96.

[10] Allen, JF, Guinn, CI, Horvitz, E. Mixed-initiative interaction. IEEE Intelligent Systems and their Applications 1999;14(5):14–23.

[11] Wongsuphasawat, K, Qu, Z, Moritz, D, Chang, R, Ouk, F, Anand, A, et al. Voyager 2: Augmenting Visual Analysis with Partial View Specifications. In: CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery; 2017, p. 2648–2659. doi:10.1145/3025453.3025768.

[12] Elzen, Svd, Wijk, JJv. Small Multiples, Large Singles: A New Approach for Visual Data Exploration. Computer Graphics Forum 2013;32(3pt2):191–200. doi:10.1111/cgf.12106.

[13] El-Assady, M, Sperrle, F, Deussen, O, Keim, D, Collins, C. Visual Analytics for Topic Model Optimization based on User-Steerable Speculative Execution. IEEE Trans on Visualization and Computer Graphics 2019;25(1):374–384. doi:10.1109/TVCG.2018.2864769.

[14] Walch, A, Schwarzler, M, Luksch, C, Eisemann, E, Gschwandtner, T. Lightguider: Guiding interactive lighting design using suggestions, provenance, and quality visualization. IEEE Trans on Visualization and Computer Graphics 2019;26(1):569–578. doi:10.1109/TVCG.2019.2934658.

[15] Endert, A, Hossain, MS, Ramakrishnan, N, North, C, Fiaux, P, Andrews, C. The human is the loop: new directions for visual analytics. Journal of Intelligent Information Systems 2014;43(3):411–435. doi:10.1007/s10844-014-0304-9.

[16] Crouser, RJ, Ottley, A, Chang, R. Balancing Human and Machine Contributions in Human Computation Systems. In: Handbook of Human Computation. Springer; 2013, p. 615–623. doi:10.1007/978-1-4614-8806-4_48.

[17] Cook, K, Cramer, N, Israel, D, Wolverton, M, Bruce, J, Burtner, R, et al. Mixed-initiative visual analytics using task-driven recommendations. In: 2015 IEEE Conference on Visual Analytics Science and Technology (VAST). 2015, p. 9–16. doi:10.1109/VAST.2015.7347625.

[18] Ceneda, D, Gschwandtner, T, Miksch, S. A Review of Guidance Approaches in Visual Data Analysis: A Multifocal Perspective. Computer Graphics Forum 2019;38(3):861–879. doi:10.1111/cgf.13730.

[19] van Wijk, J. Views on Visualization. IEEE Trans on Visualization and Computer Graphics 2006;12(4):421–432. doi:10.1109/TVCG.2006.80.

[20] Andrienko, N, Lammarsch, T, Andrienko, G, Fuchs, G, Keim, D, Miksch, S, et al. Viewing Visual Analytics as Model Building. Computer Graphics Forum 2018;37(6):275–299. doi:10.1111/cgf.13324.

[21] Federico, P, Wagner, M, Rind, A, Amor-Amorós, A, Miksch, S, Aigner, W. The Role of Explicit Knowledge: A Conceptual Model of Knowledge-Assisted Visual Analytics. In: IEEE Conf. on Visual Analytics Science and Technology. 2017, p. 92–103. doi:10.1109/VAST.2017.8585498.

[22] Sawaragi, T. Dynamical and Complex Behaviors in Human-Machine Co-Adaptive Systems. IFAC Proc Volumes 2005;38(1):94–99. doi:10.3182/20050703-6-CZ-1902.01418.

[23] Gallina, P, Bellotto, N, Di Luca, M. Progressive co-adaptation in human-machine interaction. In: Int. Conf. on Informatics in Control, Automation and Robotics; vol. 02. 2015, p. 362–368.

[24] Nikolaidis, S, Nath, S, Procaccia, A, Srinivasa, S. Game-Theoretic Modeling of Human Adaptation in Human-Robot Collaboration. In: Proc. ACM/IEEE Int. Conf. on Human-Robot Interaction. ACM; 2017, p. 323–331. doi:10.1145/2909824.3020253.

[25] Odora, RJ. Using explanation as a teaching method: How prepared are high school technology teachers in free state province, south africa? Journal of Social Sciences 2014;38(1):71–81. doi:10.1080/09718923.2014.11893238.

[26] van Liere, R, D. Mulder, J, van Wijk, JJ. Computational steering. Future Generation Computer Systems 1997;12(5):441–450. doi:10.1016/S0167-739X(96)00029-5.

[27] El-Assady, M, Sevastjanova, R, Sperrle, F, Keim, D, Collins, C. Progressive Learning of Topic Modeling Parameters: A Visual Analytics Framework. IEEE Transactions on Visualization and Computer Graphics 2018;24(1):382–391. doi:10.1109/TVCG.2017.2745080.

[28] Kim, H, Choo, J, Park, H, Endert, A. InterAxis: Steering Scatterplot Axes via Observation-Level Interaction. IEEE Transactions on Visualization and Computer Graphics 2016;22(1):131–140. doi:10.1109/TVCG.2015.2467615.

[29] Endert, A. Semantic Interaction for Visual Analytics: Inferring Analytical Reasoning for Model Steering. Synthesis Lectures on Visualization 2016;4(2):1–99. doi:10.2200/S00730ED1V01Y201608VIS007.

[30] Fiebrink, R, Gillies, M. Introduction to the Special Issue on Human-Centered Machine Learning. ACM Transactions on Interactive Intelligent Systems 2018;8(2):7:1–7:7. doi:10.1145/3205942.

[31] Kirschner, F, Paas, F, Kirschner, PA. A cognitive load approach to collaborative learning: United brains for complex tasks. Educational psychology review 2009;21(1):31–42.

[32] Schunk, DH. Learning theories an educational perspective sixth edition. Pearson; 2012.

[33] Flemisch, F, Abbink, D, Itoh, M, Pacaux-Lemoine, MP, Weßel, G. Shared control is the sharp end of cooperation: Towards a common framework of joint action, shared control and human machine cooperation. IFAC-PapersOnLine 2016;49(19):72–77. doi:10.1016/j.ifacol.2016.10.464.

[34] Felder, RM, Brent, R. The ABC's of engineering education: ABET, Bloom's taxonomy, cooperative learning, and so on. In: Proceedings American Society for Engineering Education Annual Conference & Exposition. 2004, p. 1–12.

[35] Wortman Vaughan, J, Wallach, H. A Human-Centered Agenda for Intelligible Machine Learning. In: Machines We Trust: Getting Along with Artificial Intelligence. ISBN 978-0-262-54209-8; 2021, p. 224.

[36] Wexler, J, Pushkarna, M, Bolukbasi, T, Wattenberg, M, Viégas, F, Wilson, J. The What-If Tool: Interactive Probing of Machine Learning Models. IEEE Transactions on Visualization and Computer Graphics 2019;:1–1doi:10.1109/TVCG.2019.2934619.

[37] Sperrle, F, Bernard, J, Sedlmair, M, Keim, DA, El-Assady, M. Speculative Execution for Guided Visual Analytics. In: Workshop for Machine Learning from User Interaction for Visualization and Analytics at IEEE VIS. 2018,.

[38] Burns, A, Xiong, C, Franconeri, S, Cairo, A, Mahyar, N. How to evaluate data visualizations across different levels of understanding. In: 2020 IEEE Evaluation and Beyond - Methodological Approaches for Visualization (BELIV). 2020,.

[39] Adar, E, Lee, E. Communicative Visualizations as a Learning Problem. IEEE Transactions on Visualization and Computer Graphics 2021;27(2):946–956. doi:10.1109/TVCG.2020.3030375.

[40] Saenz, M, Baigelenov, A, Hung, YH, Parsons, P. Reexamining the cognitive utility of 3D visualizations using augmented reality holograms. In: Proc. 2017 IEEE VIS Workshop on Immersive Analytics: Exploring Future Visualization and Interaction Technologies for Data Analytics. 2017,.

[41] Mahyar, N, Kim, SH, Kwon, BC. Towards a Taxonomy for Evaluating User Engagement in Information Visualization. In: Personal Visualiza-

tion : Exploring Data in Everyday Life : An IEEE VIS 2015 Workshop. 2015,URL: https://kops.uni-konstanz.de/handle/123456789/45094.
[42] Fuchs, J, Isenberg, P, Bezerianos, A, Miller, M, Keim, D. EduClust -A Visualization Application for Teaching Clustering Algorithms. In: Proc. Eurographics Conf. 2019, p. 1. URL: https://hal.inria.fr/hal-02103293.
[43] Roberts, JC, Ritsos, PD, Jackson, JR, Headleand, C. The Explanatory Visualization Framework: An Active Learning Framework for Teaching Creative Computing Using Explanatory Visualizations. IEEE Transactions on Visualization and Computer Graphics 2018;24(1):791–801. doi:10.1109/TVCG.2017.2745878; conference Name: IEEE Transactions on Visualization and Computer Graphics.
[44] El-Assady, M, Jentner, W, Kehlbeck, R, Schlegel, U, Sevastjanova, R, Sperrle, F, et al. Towards xai: Structuring the processes of explanations. 2019, p. 1–9.
[45] Miller, T. Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence 2019;267:1–38.
[46] Miller, T, Howe, P, Sonenberg, L. Explainable AI: beware of inmates running the asylum or: How I learnt to stop worrying and love the social and behavioural sciences. CoRR 2017;abs/1712.00547. arXiv:1712.00547.
[47] Wang, D, Yang, Q, Abdul, A, Lim, BY. Designing theory-driven user-centric explainable ai. In: CHI Conf. on Human Factors in Computing Systems. Association for Computing Machinery. ISBN 9781450359702; 2019, p. 1–15. doi:10.1145/3290605.3300831.
[48] Shneiderman, B. Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. International Journal of Human–Computer Interaction 2020;36(6):495–504. doi:10.1080/10447318.2020.1741118.
[49] Sacha, D, Stoffel, A, Stoffel, F, Kwon, BC, Ellis, G, Keim, DA. Knowledge Generation Model for Visual Analytics. IEEE Trans on Visualization and Computer Graphics 2014;20(12):1604–1613. doi:10.1109/TVCG.2014.2346481.
[50] Liu, J, Dwyer, T, Tack, G, Gratzl, S, Marriott, K. Supporting the Problem-Solving Loop: Designing Highly Interactive Optimisation Systems. IEEE Transactions on Visualization and Computer Graphics 2020;:1–1doi:10.1109/TVCG.2020.3030364.
[51] Karer, B, Scheler, I, Hagen, H, Leitte, H. ConceptGraph: A Formal Model for Interpretation and Reasoning During Visual Analysis. Computer Graphics Forum 2020;39(6):5–18. doi:10.1111/cgf.13899.
[52] Anderson, LW, Krathwohl, DR. A taxonomy for learning, teaching, and assessing : a revision of Bloom's taxonomy of educational objectives. Longman,; 2001. ISBN 978-0-8013-1903-7.
[53] Fekete, JD, Fisher, D, Nandi, A, Sedlmair, M. Progressive Data Analysis and Visualization (Dagstuhl Seminar 18411). Dagstuhl Reports 2019;8(10):1–40. doi:10.4230/DagRep.8.10.1.
[54] Endert, A, Ribarsky, W, Turkay, C, Wong, BW, Nabney, I, Blanco, ID, et al. The State of the Art in Integrating Machine Learning into Visual Analytics. Computer Graphics Forum 2017;36(8):458–486. doi:10.1111/cgf.13092.
[55] Liu, S, Wang, X, Liu, M, Zhu, J. Towards Better Analysis of Machine Learning Models: A Visual Analytics Perspective. Visual Informatics 2017;1(1):48–56. doi:10.1016/J.VISINF.2017.01.006.
[56] Liu, J, Dwyer, T, Marriott, K, Millar, J, Haworth, A. Understanding the Relationship Between Interactive Optimisation and Visual Analytics in the Context of Prostate Brachytherapy. IEEE Trans on Visualization and Computer Graphics 2018;24(1):319–329. doi:10.1109/TVCG.2017.2744418.
[57] Conati, C, Carenini, G, Toker, D, Lallé, S. Towards user-adaptive information visualization. In: Proc. AAAI Conf. on Artificial Intelligence. 2015, p. 4100–4106.
[58] Gotz, D, Zhou, MX. Characterizing Users' Visual Analytic Activity for Insight Provenance. Information Visualization 2009;8(1):42–55. doi:10.1057/ivs.2008.31.
[59] Wenskovitch, J, Zhou, M, Collins, C, Chang, R, Dowling, M, Endert, A, et al. Putting the "I" in Interaction: Interactive Interfaces Personalized to Individuals. IEEE Computer Graphics and Applications 2020;40(3):73–82. doi:10.1109/MCG.2020.2982465.
[60] Kindlmann, G, Scheidegger, C. An Algebraic Process for Visualization Design. IEEE Transactions on Visualization and Computer Graphics 2014;20(12):2181–2190. doi:10.1109/TVCG.2014.2346325.
[61] Shrinivasan, YB, Gotz, D, Lu, J. Connecting the dots in visual analysis. In: IEEE Symp. on Visual Analytics Science and Technology. IEEE; 2009, p. 123–130. doi:10.1109/VAST.2009.5333023.
[62] Kandel, S, Paepcke, A, Hellerstein, J, Heer, J. Wrangler: Interactive visual specification of data transformation scripts. In: Proc. Conf. Human Factors in Computing Systems (EA). ACM. ISBN 9781450302289; 2011, p. 3363. doi:10.1145/1978942.1979444.
[63] Simard, P, Amershi, S, Chickering, M, Edelman Pelton, A, Ghorashi, S, Meek, C, et al. Machine Teaching: A New Paradigm for Building Machine Learning Systems. Tech. Rep. MSR-TR-2017-26; 2017.
[64] Heer, J. Agency plus automation: Designing artificial intelligence into interactive systems. Proc National Academy of Sciences 2019;116(6):1844–1850. doi:10.1073/pnas.1807184115.
[65] Guidotti, R, Monreale, A, Ruggieri, S, Turini, F, Giannotti, F, Pedreschi, D. A Survey of Methods for Explaining Black Box Models. ACM Comput Surv 2018;51(5):93:1–93:42. doi:10.1145/3236009.
[66] Freitas, AA. Comprehensible classification models: a position paper. ACM SIGKDD Explorations Newsletter 2014;15(1):1–10. doi:10.1145/2594473.2594475.
[67] Tversky, A. Features of similarity. Psychological Review 1977;84(4):327–352. doi:10.1037/0033-295X.84.4.327.
[68] Langley, P. User Modeling in Adaptive Interface. In: UM99 User Modeling. CISM International Centre for Mechanical Sciences; Springer; 1999, p. 357–370. doi:10.1007/978-3-7091-2490-1_48.
[69] Cavallo, M, Demiralp, C. Clustrophile 2: Guided visual clustering analysis. IEEE Trans on Visualization and Computer Graphics 2018;doi:10.1109/TVCG.2018.2864477.
[70] Horvitz, E, Breese, J, Heckerman, D, Hovel, D, Rommelse, K. The lumiere project: Bayesian user modeling for inferring the goals and needs of software users. In: Conference on Uncertainty in artificial intelligence. 1998, p. 256–265.
[71] Kapoor, A, Lee, B, Tan, D, Horvitz, E. Interactive optimization for steering machine classification. In: SIGCHI Conference on Human Factors in Computing Systems. 2010, p. 1343–1352.
[72] Healey, CG, Dennis, BM. Interest driven navigation in visualization. IEEE transactions on visualization and computer graphics 2012;18(10):1744–1756. doi:10.1109/TVCG.2012.23.
[73] Monadjemi, S, Garnett, R, Ottley, A. Competing models: Inferring exploration patterns and information relevance via bayesian model selection. IEEE Transactions on Visualization and Computer Graphics 2020;.
[74] Ottley, A, Garnett, R, Wan, R. Follow The Clicks: Learning and Anticipating Mouse Interactions During Exploratory Data Analysis. Computer Graphics Forum 2019;38(3):41–52. doi:10.1111/cgf.13670.
[75] Wall, E, Arcalgud, A, Gupta, K, Jo, A. A Markov Model of Users' Interactive Behavior in Scatterplots. In: 2019 IEEE Visualization Conference (VIS). 2019, p. 81–85. doi:10.1109/VISUAL.2019.8933779.
[76] Garg, S, Ramakrishnan, IV, Mueller, K. A visual analytics approach to model learning. In: IEEE Symposium on Visual Analytics Science and Technology (VAST), 2010. IEEE; 2010, p. 67–74. doi:10.1109/VAST.2010.5652484.
[77] Behrisch, M, Korkmaz, F, Shao, L, Schreck, T. Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. In: IEEE Conf. on Visual Analytics Science and Technology (VAST). 2014, p. 43–52. doi:10.1109/VAST.2014.7042480.
[78] Dennig, FL, Polk, T, Lin, Z, Schreck, T, Pfister, H, Behrisch, M. FDive: Learning Relevance Models Using Pattern-based Similarity Measures. In: 2019 IEEE Conference on Visual Analytics Science and Technology (VAST). 2019, p. 69–80. doi:10.1109/VAST47406.2019.8986940.
[79] Das, S, Cashman, D, Chang, R, Endert, A. Beames: Interactive multimodel steering, selection, and inspection for regression tasks. IEEE computer graphics and applications 2019;39(5):20–32.
[80] Sacha, D, Kraus, M, Bernard, J, Behrisch, M, Schreck, T, Asano, Y, et al. Somflow: Guided exploratory cluster analysis with self-organizing maps and analytic provenance. IEEE Transactions on Visualization and Computer Graphics (TVCG) 2018;24(1):120–130. doi:10.1109/TVCG.2017.2744805.
[81] Bernard, J, Zeppelzauer, M, Sedlmair, M, Aigner, W. VIAL: a unified process for visual interactive labeling. The Visual Computer 2018;34(9):1189–1207. doi:10.1007/s00371-018-1500-3.
[82] Tang, T, Li, R, Wu, X, Liu, S, Knittel, J, Koch, S, et al. Plotthread: Creating expressive storyline visualizations using reinforcement learning. IEEE Transactions on Visualization and Computer Graphics 2020;.
[83] Micallef, L, Sundin, I, Marttinen, P, Ammad-ud din, M, Peltola, T, Soare, M, et al. Interactive Elicitation of Knowledge on Feature Relevance

Improves Predictions in Small Data Sets. In: Proc. Int. Conf. on Intelligent User Interfaces. ACM; 2017, p. 547–552. doi:10.1145/3025171.3025181.

[84] Auer, P. Using confidence bounds for exploitation-exploration trade-offs. Journal of Machine Learning Research 2002;3(Nov):397–422.

[85] Reda, K, Johnson, AE, Papka, ME, Leigh, J. Modeling and evaluating user behavior in exploratory visual analysis. Information Visualization 2016;15(4):325–339.

[86] Sperrle, F, Sevastjanova, R, Kehlbeck, R, El-Assady, M. VIANA: Visual Interactive Annotation of Argumentation. In: IEEE Conference on Visual Analytics Science and Technology (VAST). 2019, p. 12.

[87] Shao, L, Silva, N, Eggeling, E, Schreck, T. Visual exploration of large scatter plot matrices by pattern recommendation based on eye tracking. In: Workshop on Exploratory Search and Interactive Data Analytics. ACM; 2017, p. 9–16. doi:10.1145/3038462.3038463.

[88] Silva, N, Schreck, T, Veas, E, Sabol, V, Eggeling, E, Fellner, DW. Leveraging eye-gaze and time-series features to predict user interests and build a recommendation model for visual analysis. In: Proc. ACM Symp. on Eye Tracking Research & Applications. ACM. ISBN 9781450357067; 2018, p. 1–9. doi:10.1145/3204493.3204546.

[89] A. Leite, R, Gschwandtner, T, Miksch, S, Gstrein, E, Kuntner, J. Neva: Visual analytics to identify fraudulent networks. In: Computer Graphics Forum; vol. 39. Wiley Online Library; 2020, p. 344–359.

[90] Endert, A, Fiaux, P, North, C. Semantic interaction for visual text analytics. In: Proc. SIGCHI Conf. Human Factors in Computing Systems. ACM. ISBN 9781450310154; 2012, p. 473. doi:10.1145/2207676.2207741.

[91] Wall, E, Das, S, Chawla, R, Kalidindi, B, Brown, ET, Endert, A. Podium: Ranking Data Using Mixed-Initiative Visual Analytics. IEEE Trans on Visualization and Computer Graphics 2018;24(1):288–297. doi:10.1109/TVCG.2017.2745078.

[92] Gao, T, Dontcheva, M, Adar, E, Liu, Z, Karahalios, KG. Datatone: Managing ambiguity in natural language interfaces for data visualization. In: Proc. Symp. on User Interface Software & Technology. ACM. ISBN 9781450337793; 2015, p. 489–500. doi:10.1145/2807442.2807478.

[93] Fishburn, P. Additive utilities with incomplete product set: Applications to priorities and assignments baltimore. MD: ORSA Publication 1967;.

[94] Du, F, Plaisant, C, Spring, N, Shneiderman, B. Eventaction: Visual analytics for temporal event sequence recommendation. In: IEEE Conference on Visual Analytics Science and Technology (VAST). IEEE; 2016, p. 61–70.

[95] Jing, Y, Pavlović, V, Rehg, JM. Boosted bayesian network classifiers. Machine Learning 2008;73(2):155–184.

[96] Sperrle, F, El-Assady, M, Guo, G, Borgo, R, Chau, DH, Endert, A, et al. A Survey of Human-Centered Evaluations in Human-Centered Machine Learning. Computer Graphics Forum 2021;doi:10.1111/cgf.14329.

[97] Satyanarayan, A, Moritz, D, Wongsuphasawat, K, Heer, J. Vega-Lite: A Grammar of Interactive Graphics. IEEE Transactions on Visualization and Computer Graphics 2017;23(1):341–350. doi:10.1109/TVCG.2016.2599030.