

Why Visualize?

Arguments for Visual Support in Decision-Making

Dirk Streeb
Universität Konstanz

Mennatallah El-Assady
Universität Konstanz

Daniel A. Keim
Universität Konstanz

Min Chen
University of Oxford

Abstract—Just like the numerous applications of visualization, there are plenty of theoretical arguments for why visualization can aid in knowledge generation and communication. Meanwhile, to date these arguments have been presented independently, which hinders the exploitation of underlying properties of visualization in practice. We present a network of arguments that explicates not only how different arguments build on each other, but also which arguments demand for trade-offs in the design of visualizations. Furthermore, we call for the extension of our network to the application of visualization in decision-making, as well as visual analytics. We highlight the practical applicability of our theoretical approach with examples, including one from text analysis.

■ **DECISION-MAKING** is one of the most prominent areas of application for visualizations (see [1], for a review), next to knowledge generation and communication. Meanwhile, from a theoretic point of view there is no comprehensive and convincing explanation for the successes and failures of visualizations, yet [2]. We started to construct a network of more than 100 arguments on the merits of visualization [3].

Our network of arguments [3] provides a structural foundation for the interdisciplinary research on visualization. It combines different argumentative standpoints and relates the properties of visualization to the visualization process, as well as to their dependencies on tasks, users, and data. It also includes some limitations of visualization and highlights needs for trade-offs between opposing properties, for instance, be-




Figure 1: Small section of our network of arguments showing the foundations of potential benefits of visual patterns. We identified three groups of arguments that support the commonly expressed argument on visual patterns. Adapted from Streeb et al. [3, Fig. 5].

tween the flexible usage of visualizations and their specific tailoring to particular tasks. The plethora of properties and task demands describes the difficulties in measuring the quality of visualizations and proposing suitable views.

To give one example, let us consider the prominent argument that viewers can identify visual patterns in visualizations, which resemble patterns in the data. This argument in favor of visualization is connected to several other arguments. In particular, we explicate that the argument on visual patterns builds on arguments from at least three groups, as depicted in Figure 1. First, perception plays a role not only by driving attention towards some visual structures, but also by performing visual inferences. Second, the visual pattern can only be interpreted in case it provides precise and detailed access to the underlying data. Finally, viewers learn how to interpret visual patterns by interacting with the visualization and/or similar visualizations.

Explicating the relationships between the arguments, and thus the properties of visualization that the arguments cover, provides a better theoretical understanding of the benefits and limitations of visualization. As a next step, it is crucial to take our findings from the abstract *why visualization works* to the more practical *how does visualization work*, and consequently *how to make visualizations work*. The transfer of theoretical insights will benefit the use of visualization in decision-making and visual analytics.

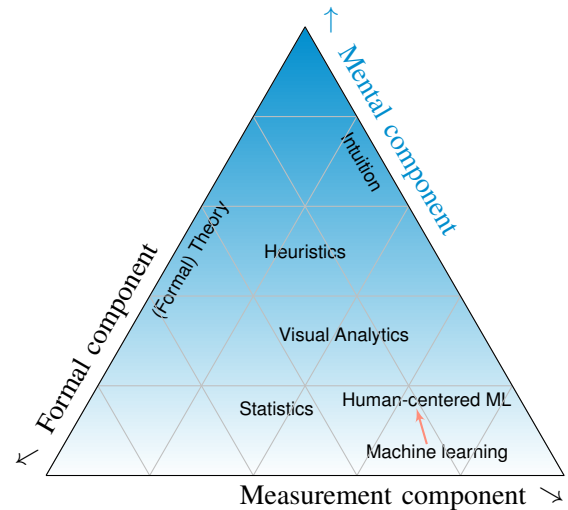


Figure 2: Decision-making strategies differ in the proportions to which the three basic components are involved. Human-centered machine learning relaxes the data fixation of machine learning.

Decision-Making Strategies

Decision-making always takes place in practical environments [4], which typically contrast the assumptions made by many mathematical models [5]. Thus, any decision-making process starts from the real world and ends by acting within the same real world, including all the costs and consequences. It does not merely connect data to the outputs of an algorithm (e.g., class labels in classification). For instance, decision-making involves different kinds of costs for measuring attribute values (e.g., time, money, side-effects) and depending on how (much) decisions are off the mark (e.g., false positives vs. false negatives).

In Figure 2, we roughly contrast a number of decision-making strategies. Strategies differ with respect to the proportions of three basic components. First, measurement can be used to quantify attributes and compile datasets, which may, for example, be used to train machine learning models. Secondly, formal models of the application domain can be used to describe the relationships between attributes mathematically, such as by laws of physics. These two components can be seen as different kinds of data. By contrast, the third component covers the not quantified aspects. This mental component includes implicit knowledge about the task and domain expertise that, for

instance, lead to the detection and interpretation of visual patterns. Here, we do not want to go into more details than having one component, which covers everything that is neither formal nor measured.

Importantly, practical decision-making strategies combine components. For instance, intuition combines cognitive processes with perceived observations and thereby balances the limitations of the mental and measurement components [4], [5]. Machine learning focuses on measured datasets, in contrast to traditional statistics that defines formal relationships between attributes in more detail based on domain-specific theories. Which (combinations of) strategies lead to appropriate and intelligent decisions depends on the application environment. For example, physicists achieved great successes in applying formal theories such as Maxwell’s equations of electromagnetism or Einstein’s theory of relativity.

The current trend to automatize machine learning processes explores the bottom (right) of the triangle and promises to relax demands on the other two components. On the one hand, it aims at compensating a potential lack of domain expertise and avoiding costs introduced by incorporating the mental component, and thus humans. On the other hand, learning relationships from measurements instead of formalizing them reduces the demand for domain-specific theories, as well as the ability to express these theories in mathematical and statistical terms.

However, the focus on measured data increases the need for high quality datasets. Though, the structure of the decision environment, structural breaks, and costly or time-consuming measurement procedures can prohibit the timely access to sufficient datasets (see also [5]). High quality measurements typically also require domain expertise for coming up with reliable measurement procedures in the first place.

Further uncertainty is introduced by the quantization/formalization of tasks, as it often does not faithfully cover all constraints. Usually a single (formal) objective function needs to be known *a priori*, which is the exact opposite of the data-driven appeal of machine learning, and demands for the formalization of domain constraints.

Needless to say, machine learning is applied, many times successfully. But even then, accessi-




Figure 3: Visualization can bring the distinct components together. Task constraints often can not be exactly formalized (e.g., the price constraint $\lesssim \$1m \implies \leq \$1m$) and measurements are rarely exact (e.g., negotiable house prices). Demands on visualizations diverge depending on the proportions of basic components involved.

bility and auditability, which may be required by contractors or law, suffer from the weak incorporation of the mental and formal components. For example, complex machine learning models may facilitate the training process, but hide the effective structures in large networks [6].

Visual Analytics

Visual analytics promotes to include humans and the mental component explicitly [7]. Visualization provides the powerful interfaces to models (e.g., from machine learning), datasets, and the general analysis process (e.g., via provenance). With our theoretical work on visualization [3], we lay a foundation for the field and guide future research. For example, we integrate theoretic work from the visualization community with arguments from psychology and other disciplines. Any human-centric approach needs to consider not only the interplay between visualizations and analysts, but also the work that analysts contribute [8], such as externalizing task constraints and embedding the analysis within the application environment. Visualization can aid analysts, for instance, by providing an external memory and a

shared structure for integrating the knowledge of experts across involved domains.

One particular strength of visualization is the potential to combine the three distinct components involved in decision-making, as sketched in [Figure 3](#). Within a single visualization the diverse entities from the three components (e.g., formal functions, data points, and mental task constraints) are mapped to a more homogeneous set of visual entities (e.g., lines, circles, and polygons). These visual entities often can be related more easily, if relevant aspects are faithfully mapped to visual features.

Let us consider the simple example of buying a house. Clearly, (I, in [Figure 3](#)) the potential buyer has a quite good mental picture of the task and the constraints on price, size, and whether the house should have a garden or a fireplace. For example, offers likely include some ridiculously expensive houses that are not feasible with the given budget. Visualizing (V) available offers can provide an overview. Machine learning (e.g., in the form of a decision tree) can help in finding a suitable house, once the potential buyer provided labels for some offers. Depicting (III) the trained model can enable the potential buyer to grasp how well it aligns with her expectations.

On their own, the three components can be valuable, but only their combination fully exploits the power of visualization. Visual interactive machine learning (IV) makes models more accessible by showing their formal and measurement components together. Visual analytics tries to exploit the analyst's domain expertise to improve models and counter limitations of data quality. For instance, (II) joining formal and mental components aids in explicating and refining task constraints, and (VI) labeling and annotation enriches datasets by externalizing domain knowledge.

We envision that the combination of all three components can be fostered. From the perspective of visual analytics there is a need to provide closer interactions with machine learning models. In complement, machine learning may benefit from the more concise alignment of intended tasks and formalized optimization criteria through visual means for evaluating models. This also demands for the exploration of the limits of visualization and resulting needs for trade-offs in practical human-machine environments, which

goes beyond our investigations to date.

Application of Visual Analytics

While there already are numerous application examples in the field of visual analytics, and more recently visual interactive machine learning [10], only lately Sacha et al. contributed an ontology for sharing a general workflow and terminology [11]. Making the iterative process human-centric enables the interactive externalization and evaluation of constraints that are, at the start, only available mentally, and often cannot be externalized faithfully in all their subtleties. Consider the following example from text analysis.

Topic modeling algorithms have become a prominent machine learning approach for automatically subdividing document corpora into thematically-related groups of documents, labeled by keywords (i.e., topic descriptors). While these models are widely applicable, they are notoriously difficult to optimize and adapt to analysts' expectations. To tackle this challenge for different user groups, the [lingvis.io](#) framework [12] offers three visual analytics techniques for refining topic models. Each of the three techniques prioritizes different components of our model.

For instance, focusing on the mental component analysts externalize their domain knowledge using [Semantic Concept Spaces](#) [13]. Here, they define high-level relations between concepts to adapt the underlying language model used for measuring word relations that form the basis of the subsequent topic modeling. The second approach [14] focuses on the formal component. Based on Layered Topic Matching [15], analysts are given a document selection and asked to compare the suitability of two topic modeling results. Using a reinforcement learning loop, the system progressively adjusts the topic modeling parameters to mirror human expectations.

Lastly, the third approach [9] emphasizes the training of the model by combining the formal and the measurement components. Here, analysts can follow along while an incremental topic model is being constructed and optimized. This approach integrates the human using different dynamics. In the following, we illustrate this last technique and how it implements our first attempt to integrate all three components.

As depicted in [Figure 4a](#), analysts start the

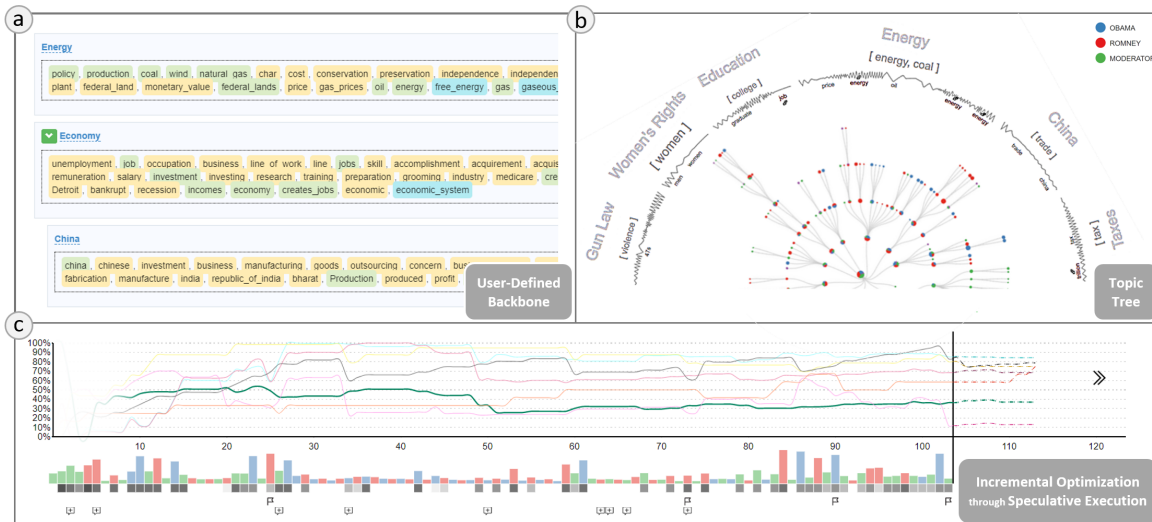


Figure 4: Interactive Topic Model Optimization using Visual Analytics. This approach [9] is based on Incremental Hierarchical Topic Modeling (IHTM) to refine a topic tree, employing *speculative execution* of different optimization strategies (based on multi-objective optimization of quality metrics).

tailored refinement process by externalizing their domain understanding in a topic backbone (i.e., a prior for priming the topic modeling algorithm). Different tasks demand for distinct backbones to faithfully reflect the diverse set of understandings of domain experts. Based on this given structure, the Incremental Hierarchical Topic Modeling (IHTM) starts constructing a topic tree that is updated on the visual interface for every measured document. Our example in Figure 4b shows a portion of a topic tree from a presidential debate between Romney and Obama. Since for analyzing such data the semantic understanding of concept relationships is central, we depict the incremental development of the topic tree. Thus, analysts are able to interrupt the processing to include their domain knowledge at any time.

Besides the refinement, the model itself utilizes a multi-objective optimization to monitor its internal quality. These objectives are various quality metrics that each capture one aspect of an optimal topic model, such as thematic coherence within a topic and separation across topics, as shown in Figure 4c. Using different trigger strategies, the visual analytics system halts the topic modeling process when the quality of the model deteriorates and activates optimizations, such as splitting large topics or merging related ones.

As potentially multiple optimizations are

beneficial (depending on the analyst’s mental model) and the optimization criteria are task-dependent (not known *a priori*), several approaches are performed in parallel using Speculative Execution (SpecEx) [16]. These speculations create sandboxes to evaluate the impact of a selected optimization on the model (projected to the future). Using such steerable preview mechanisms, the system offers options, which enable choices and allow for making actual decisions and trade-offs. Hence, this topic model optimization approach uses visual analytics to connect the formal (i.e., the hierarchical topic model), the measurement (i.e., the documents and their individual effects on the topic tree), and the mental (i.e., the prior backbone and the steering of the incremental process) components.

Making decisions and involved trade-offs demands for knowing alternatives as, for instance, generated by Speculative Execution. Machine learning should be used to come up with these options in the form of different models, instead of prescribing one model as optimal based on a single criterion that has been defined *a priori*. Visualization and visual analytics can aid in the multi-objective optimizations that are relevant to decision-making, for example, by facilitating the evaluation and comparison of models. Figure 4c shows several quality measures and their changes

throughout the modeling process. In combination with difference views for comparing topic model trees [9] the system integrates relevant information visually and enables direct interaction. As a result, our visual analytics system, while highlighting the training process and thus the formal and measurement components, also promotes the externalization of domain knowledge. This human-centric exchange of information between the machine and the human incorporates the mental component (cf. [8]).

Conclusion

The decision-making strategies in Figure 2 are options that are appropriate in some situations and inappropriate in others. We see a need to investigate the big picture of visualization within the decision-making environments it is applied in. One of the strengths of visualization is its encompassing nature regarding the combination of the different components as depicted in Figure 3. While we see this potential, we are aware of many open questions on how to combine the distinct components, for instance, in the form of sketching, mathematical diagrams and data visualization. At the same time, the application of visualization may boost creativity not only with regard to the (machine learning) models being inspected, but also with respect to how to improve the visual interfaces and visual analytics systems themselves. Our network of arguments [3] connects the beneficial properties of visualization, such that designers can navigate and integrate the theoretical aspects to consider more effectively.

Machine learning is a technological strategy for decision-making that needs to be investigated in application environments and not only as a set of purely mathematical or computational problems. Combining the fields of visualization/visual analytics and machine learning in practical contexts as visual human-centered machine learning is a major challenge for research. Extending our network on visualization towards visual analytics and visual interactive machine learning will contribute to underpin and advance the design of visual interfaces. While machine learning facilitates the exploitation of large and high-dimensional datasets, it promotes conservative one-dimensional decision-making based on past data and single optimiza-

tion criteria postulated *a priori*. Visual analytics and visual human-centered machine learning enable decision-makers to make pro-active, as well as adaptively balanced decisions based on multiple objectives evolving over time and competitive modeling options provided by a number of machine learning models.

ACKNOWLEDGMENT

We thank Hanna Schäfer. Part of this work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) within the project VALIDA/SPP-1999 RATIO (Project-ID 376714276).

REFERENCES

1. L. M. Padilla, S. H. Creem-Regehr, M. Hegarty, and J. K. Stefanucci, "Decision making with visualizations: a cognitive framework across disciplines." *Cognitive Research: Principles and Implications*, vol. 3, no. 29, 2018.
2. M. Chen, G. G. Grinstein, C. R. Johnson, J. Kennedy, and M. Tory, "Pathways for theoretical advances in visualization," *Computer Graphics and Applications*, vol. 37, no. 4, pp. 103–112, 2017.
3. D. Streeb, M. El-Assady, D. A. Keim, and M. Chen, "Why visualize? Untangling a large network of arguments," *Trans. Visualization and Computer Graphics*, vol. forthcoming, 2019.
4. G. Gigerenzer, P. M. Todd, and T. A. R. Group, Eds., *Ecological Rationality: Intelligence in the world*. Oxford University Press, 1999.
5. N. N. Taleb, *Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications*. RESEARCHERS.ONE, 2020. [Online]. Available: <https://www.researchers.one/article/2020-01-21>
6. J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," in *Int. Conf. Learning Representations*, 2019.
7. D. A. Keim, J. Kohlhammer, G. P. Ellis, and F. Mansmann, *Mastering the Information Age Solving Problems with Visual Analytics*. Eurographics Association, 2010.
8. G. K. L. Tam, V. Kothari, and M. Chen, "An analysis of machine- and human-analytics in classification," *Trans. Visualization and Computer Graphics*, vol. 23, no. 1, pp. 71–80, 2017.
9. M. El-Assady, F. Sperrle, O. Deussen, D. Keim, and C. Collins, "Visual Analytics for Topic Model Optimization based on User-Steerable Speculative Execution," *Trans. Visualization and Computer Graphics*, vol. 25, no. 1, pp. 374–384, 2019.

10. F. M. Hohman, M. Kahng, R. Pienta, and D. H. Chau, "Visual analytics in deep learning: An interrogative survey for the next frontiers." *Trans. Visualization and Computer Graphics*, vol. 25, no. 8, pp. 2674–2693, 2018.
11. D. Sacha, M. Kraus, D. A. Keim, and M. Chen, "Vis4ml: An ontology for visual analytics assisted machine learning," *Trans. Visualization and Computer Graphics*, vol. 25, no. 1, pp. 385–395, 2019.
12. M. El-Assady, W. Jentner, F. Sperrle, R. Sevastjanova, A. Hautli-Janisz, M. Butt, and D. Keim, "lingvis.io - A Linguistic Visual Analytics Framework," in *Proc. Association for Computational Linguistics, ACL System Demonstrations*. ACL, 2019.
13. M. El-Assady, R. Kehlbeck, C. Collins, D. Keim, and O. Deussen, "Semantic concept spaces: Guided topic model refinement using word-embedding projections," *Trans. Visualization and Computer Graphics*, vol. 26, no. 1, pp. 1001 – 1011, 2020.
14. M. El-Assady, R. Sevastjanova, F. Sperrle, D. Keim, and C. Collins, "Progressive learning of topic modeling parameters: A visual analytics framework," *Trans. Visualization and Computer Graphics*, vol. 24, no. 1, pp. 382–391, 2018.
15. M. El-Assady, F. Sperrle, R. Sevastjanova, M. Sedlmair, and D. Keim, "LTMA: Layered topic matching for the comparative exploration, evaluation, and refinement of topic modeling results," in *Int. Symp. Big Data Visual and Immersive Analytics*, Oct 2018, pp. 1–10.
16. F. Sperrle, J. Bernard, M. Sedlmair, D. Keim, and M. El-Assady, "Speculative execution for guided visual analytics," *arXiv preprint arXiv:1908.02627*, 2019.

Min Chen is professor of scientific visualization at Oxford University, UK, and fellow of Pembroke College. Contact him at min.chen@oerc.ox.ac.uk.

Dirk Streeb is Ph.D. candidate at the Graduate School of Decision Sciences and member of the Data Analysis and Visualization Group at the Universität Konstanz, Germany. He received his M.Sc. in Social and Economic Data Analysis from the Universität Konstanz in 2016. His research interests include theoretical aspects of information visualization and visual analytics as aids in human decision-making processes. Contact him at dirk.streeb@uni-konstanz.de.

Mennatallah El-Assady is with the Data Analysis and Visualization Group at the Universität Konstanz, Germany. Contact her at mennatallah.el-assady@uni-konstanz.de.

Daniel A. Keim is professor and head of the Data Analysis and Visualization Group at the Universität Konstanz, Germany. Contact him at keim@uni-konstanz.de.