EuroVis Workshop on Visual Analytics (2024)
M. El-Assady and H.-J. Schulz (Editors)

Extending the Visual Data Exploration Loop towards Trustworthy Machine Learning in the Healthcare Domain

D. Antweiler[1,2] and G. Fuchs[1]

[1]Fraunhofer IAIS, Sankt Augustin, Germany
[2]Fraunhofer Center for Machine Learning, Sankt Augustin, Germany

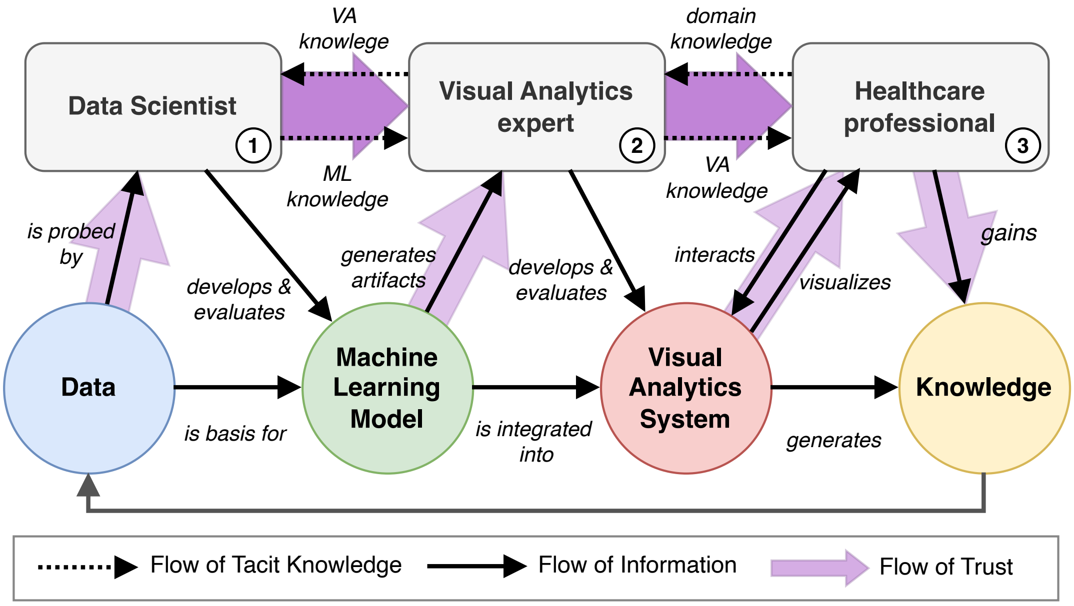

Figure 1: *Overview of the proposed visual analytics framework that fosters trust into healthcare machine learning. Circles correspond to stages from the original framework. Rounded rectangles introduce the roles of domain experts along the process. Dotted arrows represent knowledge transfer, continuous arrows reference transfer of information, while violet wide arrows indicate flow of trust. Our focus lies specifically on the interprofessional gap, in which the VA expert acts as a facilitator for the multidisciplinary team (top violet arrows).*

Abstract

Integration of machine learning (ML) systems into healthcare settings creates novel opportunities, including pattern recognition in heterogeneous medical datasets, clinical decision support as well as processes automation to save time, advance the quality of care, reduce costs and relieve healthcare staff. Challenges include opaque digital systems, curbed autonomy as well as requirements on communication, interaction and human-machine decision-making. Obstacles involve the interprofessional gap *between data scientists and healthcare professionals (HCPs) during model development as well as the* lack of trust *into ML models. Visual Analytics (VA) enables versatile interactions between users and ML models via adaptable visualizations and has been successfully deployed to improve accuracy, identify bias and increase trust. However, specifically supporting HCPs to gain trust into ML models through VA systems is not sufficiently explored. We propose an extended visual data exploration framework towards trustworthy ML in the healthcare domain for multidisciplinary teams of data scientists, VA experts and HCPs. Additionally, we apply our framework to three real-world use cases for policy development, plausibility testing and model optimization.*

CCS Concepts

*• **Applied computing** → **Health care information systems;** • **Computing methodologies** → **Machine learning;***

DOI: 10.2312/eurova.20241107

1. Introduction

Machine learning (ML) applications in the healthcare domain are spreading steadily and ever quicker. Due to increasing patient counts, amount of data, as well as number and pace of decisions being made, both the needs and benefits of applying ML techniques are greater than ever. There are distinct opportunities for ML to reduce healthcare costs, decrease disease burden and increase positive outcomes for patients, as well as satisfaction of both patients and staff. Machine learning is therefore an important part of in the ongoing digital transformation of many healthcare systems worldwide.

While many applications for ML in healthcare have been proposed, only few have been successfully deployed in real-world clinical settings. Two main obstacles are the *interprofessional gap* between data scientists and healthcare professionals during model development as well as the *lack of trust* into machine learning models [OSV21]. These gaps lead to disruptions in communication and decision-making. Visual Analytics (VA) methods enable versatile interactions between users and machine learning models. They have been successfully deployed to improve accuracy, identify bias and increase trust [BABB*21]. Central to this task is the conjunction of the users mental models (e.g. pre-existing, tacit or process knowledge) with the model's characteristics and outputs [ALA*18]. However, specifically supporting healthcare professionals with trustworthy machine learning models through Visual Analytics systems is not sufficiently explored [OSV21].

We can distill the described problem setting into the following research question and sub-questions for guidance, to which we will refer to in the remainder of this work:

Q: How can visual analytics support healthcare professionals and data scientists to build trust in ML models for decision support?
$\mathbf{SQ_1}$: How can we build trust during *complex decision-making*?
$\mathbf{SQ_2}$: How can we build trust into models based on *complex data*?
$\mathbf{SQ_3}$: How can we build trust into *complex models*?

Within the intersection of trustworthy ML and healthcare, our contributions are three-fold: (1) We identify a research gap through a survey of existing visual data analysis frameworks (Tab. 1), (2) propose an extended framework towards visual data exploration that can act as a guideline to design clinical VA systems (**Q**, Fig. 1), and (3) show the rationale of the proposed framework by contextualizing three VA tools each addressing real-world settings spanning different healthcare areas, users and contexts ($\mathbf{SQ_1} - \mathbf{SQ_3}$).

2. Related Work

2.1. Trust in Visual Analytics

Trust is a fundamental factor in how users engage with Visual Analytics systems. It can be differentiated into *cognitive* vs. *affective* trust [ESB*23], where both categories exhibit multiple *trust antecedents* relating to either the data or the visualization part of a given system. Domain experts can have high trust even in novel visual analytics systems provided they are intuitive, transparent and flexible enough to let users switch easily between analysis tasks [DLW*17]. Consequently, some studies try to measure trust as integral aspects of visualizations [PFCB23]. Others describe a continuous scale from distrust to trust, with a *limit of forgivability* and a *cooperation threshold* in between [HS20].

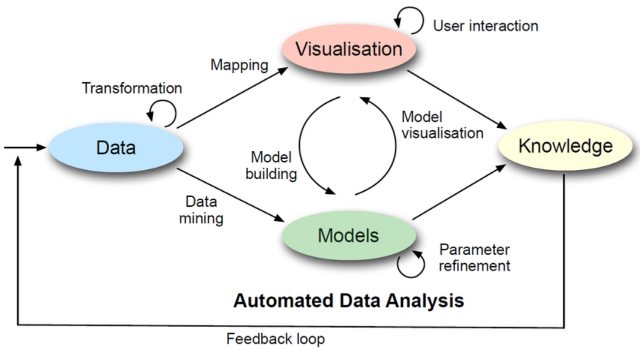

Figure 2: *Visual Data Exploration Loop consisting of data, models, visualisation and knowledge connected by their respective relationships upon which our framework builds [KMSZ09].*

Visual analytics systems can have a trust advantage in this respect over other information-based mediums especially for complex analytical reasoning tasks. For example, trust into models can be increased by visualizing their specific uncertainties [PLP*22], describing the provenance of the training datasets [RESC16], externalizing trust [EAA*23], or visually comparing test samples with examples from the train set during inference [LJHW23].

2.2. Visual Data Exploration Loop and its Extensions

Presence of the above mentioned and other trust-building measures and tools in Visual Analytics systems is well aligned with core challenges of analytical reasoning [TC05], in particular when facing ambiguous or even conflicting data. Adequately supporting timely, defensible and understandable assessments and subsequent decision making in these settings requires, explicitly or implicitly, the consolidation of structures inside data with user's tacit knowledge.

The Visual Data Exploration Loop [KMSZ09] proposes a high-level conceptual model of a system supporting these requirements, containing four stages *Data*, *Visualisation*, *Models* and *Knowledge*, with transitions between them (cf. Fig. 2). While widely recognized, Keim's model did not capture trust explicitly.

We searched the scientific database Google Scholar for publications citing the original framework as well as describing themselves as extensions or are evidently based on Keim's framework since its publication and identified 17 publications, that meet this criteria; twelve of them related to at least one of our topics of interest. (cf. Tab. 1). Certain works emphasize the human-computer interaction regarding cognitive aspects [BL21], and objectives and actions of users [RAW*16]. Specific parts of the analytical process are likewise addressed in detail, including data preprocessing [MLPM21], model behavior on high-dimensional tasks [PvSE*22], model building [ALA*18] or model specifications [SCK20]. Most closely related to our work are approaches focused on human-centered machine learning [SSZ*17], knowledge discovery [SGGB13] and knowledge generation [SSS*14] as well as uncertainty, awareness and trust [SSK*16]. A singular framework to address both trustworthy machine learning as well as the specific challenges of multidisciplinary decision-making in healthcare is notably missing.

Source	Year	Area	Trust	XAI	ML	Teams
[BAF*13]	2013	model selection in time series analysis	□	□	■	□
[SGGB13]	2013	knowledge discovery	□	□	■	□
[SSS*14]	2014	knowledge generation	◪	□	□	◪
[SSK*16]	2016	uncertainty, awareness and trust	■	□	◪	□
[SSZ*17]	2017	human-centered machine learning	◪	■	■	◪
[ALA*18]	2018	model building	◪	□	◪	□
[SSB*19]	2019	gameful design concepts	□	□	□	■
[SKKC19]	2019	Visual Analytics assisted ML	□	□	■	□
[SCK20]	2020	model specifications	■	■	■	□
[MLPM21]	2021	data preprocessing and profiling	■	□	■	□
[PvSE*22]	2022	model behavior on high-dimensional tasks	□	■	■	□
[GCMH23]	2023	longitudinal clinical studies	□	□	■	□

Table 1: *Extensions to Keim's Visual Data Exploration Loop with their specific areas and whether they address the visual analytics aspects in full (■), partly (◪) or not at all (□). XAI: Explainable Artificial Intelligence, ML: Machine Learning, Teams: Multi-disciplinary teams.*

3. Framework

We exhibit our framework by introducing relevant stakeholders and their relationships as well as describing identified combined requirements with respect to the overall flow of trust between these stakeholders.

3.1. Stakeholders

Our framework makes an addition to the originally proposed framework in terms of stakeholders. We augment the visual data exploration loop with three protagonist groups present in most use cases (cf. Fig. 1):

(1) **Data Scientist**: Experts in data handling, pre-processing, selection, training and evaluation of machine learning models are tasked to probe the data to explore relevant patterns and develop a mathematical description of them. Data scientists almost always start out without explicit knowledge about the dataset's provenience, including data collection process, quality control gates or original intended cause. This complicates data pre-processing such as cleaning, imputation or transformation. Still, these steps are fundamental to the trust process of data scientists and lay the foundation for model architecture selection and evaluation processes. Along the way, these AI model builders execute multiple explicit and implicit *sanity-checks* [BSN*23]. Data Scientists can be assisted in their analysis by means of *datasheets* for datasets, which contains information about motivation and process of data collection, composition, preprocessing, distribution and maintenance [GMV*18].

(2) **Visual Analytics expert**: The design and development of a visual analytics system calls for specialists with knowledge about data visualization, user experience, human perception as well as interaction. The task of the VA expert lies in the creation of single or multiple visual interfaces containing a layout of views showcasing aspects of the data or results derived from it. This can be achieved through interactive elements such as responsive graphs or other visual elements. It is necessary, that the VA expert understands the nuanced relationships between the input data and the end-user's decision-process and can support the identification of patterns using off-the-shelf or custom visual analytics tools.

(3) **Healthcare professional**: Medical practitioners are the end users in most scenarios of the decision support system, so the system's design needs to cater to them and their workflows. This stakeholder group can include clinicians such as nursing staff, doctors, surgeons or assistants. Their primary focus lies on optimal patient treatment under the constraint of limited resources. These resources could be time, medical supplies, availability of medical devices or access to medication. When dealing with machine learning systems, their most significant need is confirmation of their experience and present appraisal, as well as linking of knowledge from multiple sources [BSN*23].

3.2. Requirements

VA systems in healthcare need to adequately address the differences in task focii, prior knowledge, and preferences of these three principal stakeholder groups. Reflection of previous works [ASR*22, AF22, AFG23] allows to identify a set of combined requirements for trustworthy VA in the healthcare domain. They were gathered during structured interviews and workshops with HCPs. These can be divided into *accessibility*, *complexity reduction* and *technical* requirements. Sec. 4 reviews how these general requirements are addressed within the context of specific use cases.

Accessibility

R_1: Use of widely familiar visualizations that are accessible both to data scientists and healthcare professionals

R_2: Use of suitable terms and definitions for both domain experts and model developers

R_3: User guidance to assist them through the analysis process, and to provide recommendations, explanations, or insights, based on the user's goals, preferences, or behavior

Complexity Reduction

R_4: Linked multiple view interface integrating all task-related aspects according to the information seeking mantra [Shn96]

R_5: Data reduction due to excessive number of features to focus on the most relevant aspects [Shn96]

R_6: Reducing complexity of employed model architecture

Proceedings published by Eurographics - The European Association for Computer Graphics.

Security & Scalability

- $\mathbf{R_7}$: Computation and visualization is done locally to reduce risks when working with privacy-sensitive data from patients
- $\mathbf{R_8}$: Scalable visualizations for up to millions of data points that circumvent overplotting
- $\mathbf{R_9}$: Fast results for near-real-time analysis to prevent slowdown of medical processes

Although other domains exhibit similar combinations of sensitive data and high risk decisions, healthcare is somewhat special due to (1) the complexity of stakeholders (HCPs, patients) with different backgrounds and intricate relationships, (2) the heterogeneous datasets covering both medical and non-medical information as well as that (3) trustworthiness of employed ML models is not simply nice-to-have but often required by law. We are ready to describe the flow of trust between the stakeholders of our framework.

3.3. Flow of Trust

We describe the information flow between these groups of stakeholders in detail to identify necessary routes of trust and argue for a systematic approach of trust building along the visual data exploration loop (cf. Fig. 1): As trust aspects regarding data, ML model and VA system have already been adequately studied (Fig. 1, bottom row of violet arrows), our focus lies specifically on the interprofessional gap, in which the VA expert acts as a facilitator for the multidisciplinary team related to cognitive trust (cf. [ESB*23]).

1 → 2 ($\mathbf{FT}_{12}$): The data scientist should communicate essential aspects of the data handling and modeling phase to the VA expert. This process starts with a detailed data understanding. Significant facets relevant to subsequent tasks include data transformations, the selection process of the model architecture, the coverage of probed hyperparameter combinations, quality metrics, uncertainties of the trained model as well as how edge cases are handled. This communication can be supported by the usage of *model cards*, containing relevant information about intended use, metrics, evaluation, training, ethical considerations and caveats [MWZ*18].

2 → 3 ($\mathbf{FT}_{23}$): Visual Analytics experts need to explicitly convey their course of action to end-users, i.e. healthcare professionals. In doing so, they should address the aspects of interaction with the visualizations, limitations of chosen views, as well as characteristics of the complexity reduction employed. While face-to-face tutorials should be preferred, user guides or visual guidance elements can support the training curve as well.

3 → 2 ($\mathbf{FT}_{32}$): HCPs such as physicians, nurses or public health experts carry substantial experience and knowledge about their individual profession and associated workflows. Regarding any decision support system put into practice, their main task is to differentiate between signal and noise within the system's results, i.e., to put identified patterns into clinical context and reflect mismatches between data analysis and reality back to VA experts.

2 → 1 ($\mathbf{FT}_{21}$): Experts for visual analytics on the one hand are required to procure the data analysis steps in general, whether data visualization principles are concerned. This includes the compliance with effectiveness, appropriateness, and expressiveness to prevent overplotting and other user obstacles. On the other hand, they need to mediate feedback from end users, in our case healthcare professionals, back to the data scientists.

4. Use Cases

In this section we review three previous works that each primarily address one of the three research questions $\mathbf{SQ_1} - \mathbf{SQ_3}$, respectively. In doing so, each outlines a case-specific instantiation of the proposed flow of trust framework, in particular, which flows and requirements $\mathbf{R_1} - \mathbf{R_9}$ are most relevant and thus chiefly informed design decisions of the VA systems presented in each paper.

4.1. Visual Analytics for Machine Learning-based Healthcare Policy Development ($\mathbf{SQ_1}$) [ASR*22]

A major challenge for departments of public health (DPHs) in dealing with the COVID-19 pandemic has been tracing contacts in exponentially growing SARS-CoV-2 infection clusters. Prevention of further disease spread requires a comprehensive registration of individuals to clusters. Due to the high number of infections with unknown origin, the healthcare analysts need to identify connected cases and clusters through accumulated epidemiological knowledge and the infection metadata in their database (R_4). [ASR*22] contributes a VA framework to identify, assess and visualize clusters in contact tracing networks (R_1). It calculates and visualizes possible missing infection routes inside the network and supports the analysis of time-dependent events that led to the spread of the virus (R_9). An essential aspect is the display of model uncertainties (FT_{12}) and detailed legends for visual elements (FT_{23}). Additionally, it demonstrates how graph-based machine learning methods can be used to find missing links between infection clusters and thus support the mission to get a comprehensive view on infection events (R_3). The addition of a visualization component displaying the spatial extent of the infection chains further improves traceability (R_8). The design of the system and the underlying data transformations mimic the traditional epidemiological analysis approaches (FT_{32}, FT_{21}). The proposed system has been developed as a responsive web-app (R_7) and positively evaluated in collaboration with DPHs.

4.2. Visual Analytics for Plausibility Testing in Healthcare Machine Learning ($\mathbf{SQ_2}$) [AF22]

Deteriorating conditions in hospital patients are a major factor in clinical patient mortality. Currently, timely detection is based on clinical experience, expertise, and attention. However, healthcare trends towards larger patient cohorts, more data, and the desire for better and more personalized care are pushing the existing, simple scoring systems to their limits. Data-driven approaches can extract decision rules from available medical coding data, which offer good interpretability and thus are key for successful adoption in practice (R_9). The proposed visual analytics system supports healthcare professionals in inspecting and enhancing such rule-based classifiers (R_1, R_2, R_6) by visualizing similarities and differences between rules (R_5), as well as contextualizing their feature distribution within hierarchical code structures (R_4). It further provides means to modify rules to match existing medical knowledge (R_3). The system was developed iteratively in close collaboration with medical professionals i.e., with an emphasis on the flow of trust between the VA expert and the healthcare professional (FT_{23}, FT_{32}).

4.3. Visual Analytics for Transformer Model Optimization in Healthcare ($\mathbf{SQ_3}$) [AFG23]

Detecting deteriorating conditions early allows for medical intervention in hospital patients. As patient data often comes in the form of structured temporal data, many approaches rely on Large Language Models such as BERT. Interpreting these models poses a significant challenge, making trust difficult to establish. This work created a visual analytics system for examining, comparing, and elucidating pre-trained transformer models used in clinical outcome prediction tasks with the overarching goal of facilitating trust building. The use case is developed on the basis of a large hospital patient dataset and prediction tasks for acute kidney injury and heart failure. A detailed dataset exploration module is integrated to facilitate trust into the training data (FT_{12}). The system comprises well-known visualization types such as histograms, line or sankey charts (R_1), detailed labeling and legends (R_2) as well as guidance through the analysis process via a suggested examination cycle and preference options (R_3, FT_{23}). A filtered and linked two-part view is employed at a suitable point to reduce complexity and focus the user's attention (R_4, R_5). This view utilizes a medical code hierarchy suitable to the clinical use case (FT_{32}). While the multi-task transformer model architecture itself is of high complexity, the system supports model exploration by visualizing pre-training strategies and training loss convergence (R_6). The tool is engineered as a lightweight locally-deployable web-application with quick responses to user inputs ($R_7 - R_9$). Discussion with HCPs confirms that such a system can lead to a faster decision process and improved modeling results.

5. Discussion

We recapitulate how our framework can guide the design and development of a trustworthy visual analytics system in healthcare ML.

Challenges and requirements along the visual data exploration loop can be divided into *computer-centered* and *human-centered*. On the *computer side*, healthcare professionals need to deal with larger patient cohorts and more diverse data as well as the need for higher personalization, taking into account a wide variety of information about an individual. On the *human side*, there is the need for understanding and trust when ML models are used by humans which is even being demanded by law. Increased trust facilitates regulatory approvals and reduces the burden of decisions. The interdisciplinary nature of the problem adds complexity as VA experts often act as intermediaries between data scientists and healthcare professionals. Our survey reveals a gap in the literature, which is missing a comprehensive framework, that adequately tackles the dual requirements of trustworthy machine learning and the unique complexities of multi-disciplinary decision-making in healthcare (cf. Tab. 1).

Trust during complex decision-making can be approached via *retracing* the current workflow to pinpoint moments in time for leveraging visual analytics and ML assistance ($\mathbf{SQ_1}$). These junctures may include instances where decision-making is predominantly reliant on intuition and experience, despite the presence of data that may not be readily utilizable in its current format. In order to establish credibility in models derived from intricate datasets ($\mathbf{SQ_2}$), it is imperative to reduce complexity, prioritize salient factors, and incorporate domain-specific terminology pertinent to the task at hand. As healthcare data is often heterogenous, of varying quality and show only a small part of reality, both ML models and VA systems are forced to adapt appropriately. This means to highlight limitations, counteract biases as well as enable versatile interactions. We can facilitate trust into complex models ($\mathbf{SQ_3}$), by making models both *explainable* and *explorable* for the end user. The former can be achieved either through interpretable models as well as through post-hoc explanations of entire models or individual predictions. The latter can be achieved through baseline comparisons and the visualization of model architecture, uncertainty, and robustness. We must add that the loop can contain additional connections not highlighted here, including a direct exchange between HCPs and Data Scientists as well as data analysis done by the VA expert.

Regarding our main research question (**Q**) we argue that data, models and visualization must not be considered individually, but as part of the real-world process including users, conditions and dependencies. VA systems need to cater to both data scientists as well as healthcare professionals. This interdisciplinary team need to create a common vocabulary to discuss the intricacies of clinical use cases. Experts in VA should realize their special intermediate role between those two groups and prepare visual elements accordingly.

Our framework is motivated by numerous projects with hospitals, public health offices and medical practices. Nevertheless, healthcare contains far more protagonists such as outpatient care, pharmacies and insurances with highly-specific processes, data pools and use cases. We would like to integrate their perspective and extend our VA framework in the future. Additionally, our framework is limited by considering *models* as fixed artifacts, which can be probed for interpretability by the user. Current developments with generative AI indicate that AI soon will be able to act as an agent itself and explain/self-explore its inputs, outputs and mode of operation.

6. Conclusion and Future Work

In this work, we presented an extension to the visual data exploration loop towards trustworthy machine learning in the healthcare domain. Our main contributions comprise the identification of a research gap through an in-depth survey, the introduction of stakeholders into the framework as well as illuminating the flow of trust along the trust-related transitions between them. The elements of the framework can be used as guiding principles for the development of visual healthcare analytics systems. We referenced three real-world use cases from different healthcare areas, in which visual analytics actively supported trust building into complex machine learning models, as confirmed by healthcare professionals.

For future work, we plan to integrate *trust quality gates* into our framework to further develop the process-related perspective. These quality gates arise along the visual data exploration loop and should be addressed through interactive visual elements. Examples include the visual communication of systematic model weaknesses, assumptions made during the visualization process and the detection of concept or model drift. Additionally, we would like to discuss optimal points in the visual data exploration process to fuse the user's individual pre-existing knowledge and expertise. The stakeholders hold different experiences and use different vocabulary. The challenge is to create a common body of knowledge as a basis for joint decision-making in high stakes situations.

References

[AF22] ANTWEILER D., FUCHS G.: Visualizing rule-based classifiers for clinical risk prognosis. In *2022 IEEE (VIS)* (Oct. 2022), IEEE. doi:10.1109/vis54862.2022.00020. 3, 4

[AFG23] ANTWEILER D., FUCHS G., GALLUSSER F.: Multi-Task Transformer Visualization to Build Trust for Clinical Outcome Prediction. In *VAHC 2023 (VIS)* (2023), IEEE. 3, 5

[ALA*18] ANDRIENKO N., LAMMARSCH T., ANDRIENKO G., FUCHS G., KEIM D., MIKSCH S., RIND A.: Viewing visual analytics as model building. *CGF 37*, 6 (2018). doi:10.1111/cgf.13324. 2, 3

[ASR*22] ANTWEILER D., SESSLER D., ROSSKNECHT M., ABB B., GINZEL S., KOHLHAMMER J.: Uncovering chains of infections through spatio-temporal and visual analysis of COVID-19 contact traces. *C & G 106* (Aug. 2022), 1–8. doi:10.1016/j.cag.2022.05.013. 3, 4

[BABB*21] BEAUXIS-AUSSALET E., BEHRISCH M., BORGO R., CHAU D. H., COLLINS C., EBERT D., EL-ASSADY M., ENDERT A., KEIM D. A., KOHLHAMMER J., OELKE D., PELTONEN J., RIVEIRO M., SCHRECK T., STROBELT H., VAN WIJK J. J.: The role of interactive visualization in fostering trust in AI. *IEEE Computer Graphics and Applications 41*, 6 (Nov. 2021), 7–12. doi:10.1109/mcg.2021.3107875. 2

[BAF*13] BOGL M., AIGNER W., FILZMOSER P., LAMMARSCH T., MIKSCH S., RIND A.: Visual analytics for model selection in time series analysis. *IEEE TVCG 19*, 12 (2013), 2237–2246. doi:10.1109/tvcg.2013.222. 3

[BL21] BOUAZIZI S., LTIFI H.: *Improved Visual Analytic Process Under Cognitive Aspects*. Springer International Publishing, 2021, p. 494–506. doi:10.1007/978-3-030-75100-5_43. 2

[BSN*23] BORYS K., SCHMITT Y. A., NAUTA M., SEIFERT C., KRÄMER N., FRIEDRICH C. M., NENSA F.: Explainable AI in medical imaging: An overview for clinical practitioners – saliency-based XAI approaches. *European Journal of Radiology 162* (May 2023), 110787. doi:10.1016/j.ejrad.2023.110787. 3

[DLW*17] DASGUPTA A., LEE J.-Y., WILSON R., LAFRANCE R. A., CRAMER N., COOK K., PAYNE S.: Familiarity vs trust: A comparative study of domain scientists' trust in visual analytics and conventional analysis methods. *IEEE TVCG 23*, 1 (Jan. 2017), 271–280. doi:10.1109/tvcg.2016.2598544. 2

[EAA*23] ELZEN S. V. D., ANDRIENKO G., ANDRIENKO N., FISHER B. D., MARTINS R. M., PELTONEN J., TELEA A. C., VERLEYSEN M.: The flow of trust: A visualization framework to externalize, explore, and explain trust in ml applications. *IEEE CG&A 43*, 2 (Mar. 2023), 78–88. doi:10.1109/mcg.2023.3237286. 2

[ESB*23] ELHAMDADI H., STEFKOVICS A., BEYER J., MOERTH E., BEARFIELD C. X., NOBRE C.: Vistrust: a multidimensional framework and empirical study of trust in data visualizations, 2023. arXiv:2309.16915. 2, 4

[GCMH23] GÓMEZ D. A., CHARPAK N., MONTEALEGRE A., HERNÁNDEZ J. T.: Vals: Supporting visual data analysis in longitudinal clinical studies. *IEEE Access 11* (2023), 28820–28830. doi:10.1109/access.2023.3259972. 3

[GMV*18] GEBRU T., MORGENSTERN J., VECCHIONE B., VAUGHAN J. W., WALLACH H., DAUMÉ H., CRAWFORD K.: Datasheets for datasets, 2018. doi:10.48550/ARXIV.1803.09010. 3

[HS20] HAN W., SCHULZ H.-J.: Beyond trust building — calibrating trust in visual analytics. In *2020 IEEE TREX* (Oct. 2020), IEEE. doi:10.1109/trex51495.2020.00006. 2

[KMSZ09] KEIM D. A., MANSMANN F., STOFFEL A., ZIEGLER H.: Visual analytics. In *Encyclopedia of Database Systems*. Springer US, 2009, pp. 3341–3346. doi:10.1007/978-0-387-39940-9_1122. 2

[LJHW23] LI H. X., JORGENSEN S., HOLODNAK J., WOLLABER A. B.: Scatteruq: Interactive uncertainty visualizations for multiclass deep learning problems. In *2023 IEEE VIS* (Oct. 2023), IEEE. doi:10.1109/vis54172.2023.00058. 2

[MLPM21] MILANI A. M. P., LOGES L. A., PAULOVICH F. V., MANSSOUR I. H.: Prava: Preprocessing profiling approach for visual analytics. *Information Visualization 20*, 2–3 (July 2021), 101–122. doi:10.1177/14738716211021591. 2, 3

[MWZ*18] MITCHELL M., WU S., ZALDIVAR A., BARNES P., VASSERMAN L., HUTCHINSON B., SPITZER E., RAJI I. D., GEBRU T.: Model cards for model reporting, 2018. doi:10.48550/ARXIV.1810.03993. 4

[OSV21] OOGE J., STIGLIC G., VERBERT K.: Explaining artificial intelligence with visual analytics in healthcare. *WIREs Data Mining and KD 12*, 1 (Nov. 2021). doi:10.1002/widm.1427. 2

[PFCB23] PADILLA L., FYGENSON R., CASTRO S. C., BERTINI E.: Multiple forecast visualizations (mfvs): Trade-offs in trust and performance in multiple covid-19 forecast visualizations. *IEEE TVCG 29*, 1 (2023), 12–22. doi:10.1109/TVCG.2022.3209457. 2

[PLP*22] PANAGIOTIDOU G., LAMQADDAM H., POBLOME J., BROSENS K., VERBERT K., MOERE A. V.: Communicating uncertainty in digital humanities visualization research. *IEEE TVCG* (2022), 1–11. doi:10.1109/tvcg.2022.3209436. 2

[PvSE*22] PRASAD V., VAN SLOUN R. J. G., ELZEN S. V. D., VILANOVA A., PEZZOTTI N.: The transform-and-perform framework: Explainable deep learning beyond classification. *IEEE TVCG* (2022), 1–14. doi:10.1109/TVCG.2022.3219248. 2, 3

[RAW*16] RIND A., AIGNER W., WAGNER M., MIKSCH S., LAMMARSCH T.: Task cube: A three-dimensional conceptual space of user tasks in visualization design and evaluation. *InfoVis 15*, 4 (July 2016), 288–300. doi:10.1177/1473871615621602. 2

[RESC16] RAGAN E. D., ENDERT A., SANYAL J., CHEN J.: Characterizing provenance in visualization and data analysis: An organizational framework of provenance types and purposes. *IEEE TVCG 22*, 1 (Jan. 2016), 31–40. doi:10.1109/tvcg.2015.2467551. 2

[SCK20] SCHLEGEL U., CAKMAK E., KEIM D. A.: ModelSpeX: Model Specification Using Explainable Artificial Intelligence Methods. In *ML Methods in Vis. for Big Data* (2020), Archambault D., Nabney I., Peltonen J., (Eds.), EG. doi:10.2312/mlvis.20201100. 2, 3

[SGGB13] STAHL F., GABRYS B., GABER M. M., BERENDSEN M.: An overview of interactive visual data mining techniques for knowledge discovery. *WIREs Data Mining and Knowledge Discovery 3*, 4 (July 2013), 239–256. doi:10.1002/widm.1093. 2, 3

[Shn96] SHNEIDERMAN B.: The eyes have it: a task by data type taxonomy for information visualizations. In *IEEE Symp. on Vis. Languages* (1996), IEEE Comput. Soc. Press. doi:10.1109/vl.1996.545307. 3

[SKKC19] SACHA D., KRAUS M., KEIM D. A., CHEN M.: Vis4ml: An ontology for visual analytics assisted machine learning. *IEEE TVCG 25*, 1 (2019), 385–395. doi:10.1109/TVCG.2018.2864838. 3

[SSB*19] SEVASTJANOVA R., SCHÄFER H., BERNARD J., KEIM D., EL-ASSADY M.: Shall we play? – extending the visual analytics design space through gameful design concepts. In *2019 IEEE MLUI Workshop* (Oct. 2019), IEEE. doi:10.1109/mlui52769.2019.10075563. 3

[SSK*16] SACHA D., SENARATNE H., KWON B. C., ELLIS G., KEIM D. A.: The role of uncertainty, awareness, and trust in VA. *IEEE TVCG 22*, 1 (2016). doi:10.1109/tvcg.2015.2467591. 2, 3

[SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G., KEIM D. A.: Knowledge generation model for VA. *IEEE TVCG 20*, 12 (2014), 1604–1613. doi:10.1109/tvcg.2014.2346481. 2, 3

[SSZ*17] SACHA D., SEDLMAIR M., ZHANG L., LEE J. A., PELTONEN J., WEISKOPF D., NORTH S. C., KEIM D. A.: What you see is what you can change: Human-centered machine learning by interactive visualization. *Neurocomputing 268* (Dec. 2017), 164–175. doi:10.1016/j.neucom.2017.01.105. 2, 3

[TC05] THOMAS J. J., COOK K. A.: *Illuminating the path: The Research and Development Agenda for Visual Analytics*. IEEE Com. Soc., 2005. 2