






The State of the Art in Enhancing Trust in Machine Learning Models with the Use of Visualizations

A. Chatzimpampas¹, R. M. Martins¹, I. Jusufi¹, K. Kucher¹, F. Rossi², and A. Kerren¹

¹Department of Computer Science and Media Technology, Linnaeus University, Sweden

²Ceremade, Université Paris Dauphine, PSL University, France

Abstract

Machine learning (ML) models are nowadays used in complex applications in various domains, such as medicine, bioinformatics, and other sciences. Due to their black box nature, however, it may sometimes be hard to understand and trust the results they provide. This has increased the demand for reliable visualization tools related to enhancing trust in ML models, which has become a prominent topic of research in the visualization community over the past decades. To provide an overview and present the frontiers of current research on the topic, we present a State-of-the-Art Report (STAR) on enhancing trust in ML models with the use of interactive visualization. We define and describe the background of the topic, introduce a categorization for visualization techniques that aim to accomplish this goal, and discuss insights and opportunities for future research directions. Among our contributions is a categorization of trust against different facets of interactive ML, expanded and improved from previous research. Our results are investigated from different analytical perspectives: (a) providing a statistical overview, (b) summarizing key findings, (c) performing topic analyses, and (d) exploring the data sets used in the individual papers, all with the support of an interactive web-based survey browser. We intend this survey to be beneficial for visualization researchers whose interests involve making ML models more trustworthy, as well as researchers and practitioners from other disciplines in their search for effective visualization techniques suitable for solving their tasks with confidence and conveying meaning to their data.

Keywords: trustworthy machine learning, visualization, interpretable machine learning, explainable machine learning

ACM CCS: • Information systems → Trust; • Human-centered computing → Visual analytics; • Human-centered computing → Information visualization; • Human-centered computing → Visualization systems and tools; • Machine learning → Supervised learning; • Machine learning → Unsupervised learning; • Machine learning → Semi-supervised learning; • Machine learning → Reinforcement learning

1. Introduction

Trust in machine learning (ML) models is one of the greatest challenges in real-life applications of ML [TAC*20]. ML models are now commonplace in many research and application domains, and they are frequently used in scenarios of complex and critical decision-making [NGDM*19, PWJ06, TKK18]. Medicine, for example, is one of the fields where the use of ML might offer potential improvements and solutions to many difficult problems [KKS*19, SGSG19, SKK*19]. A significant challenge that remains, however, is how *trustworthy* are the ML models that are being used in these disciplines. Rudin and Ustun [RU18], for example, emphasize the importance of trust for ML models in healthcare and criminal justice, since they play a significant role in making decisions regarding human lives. It is not uncommon to observe that domain experts may not rely on ML models if they do not understand how they work [JSO19].

The impact of this problem can already be observed in recent works, such as the program “Explainable AI (XAI)” founded by DARPA (Defense Advanced Research Projects Agency) [Dar20] and described by Krause et al. [KDS*17]. This initiative is only one of the various projects that suggest further research into the field of XAI, which—to a certain extent—addresses challenges related to trust. The XAI program in its two main motivational points mentions specifically that “producing more explainable models, while maintaining a high level of learning performance” and “enabling human users to understand, appropriately trust, and effectively manage the emerging generation of AI” are both key actions for the future development in numerous domains that use ML. Understanding and trusting ML models is also arguably mandatory under the General Data Protection Regulation (GDPR) [EC16] as part of the “right to be informed” principle: data controllers must provide meaningful information about the logic involved in automated decisions [Art18]. Individuals have also the right not to be subject

to a decision based solely on automated processing: enabling subjects of ML algorithms to trust their decision is probably the easiest way to reduce the objection to such automated decisions.

In reaction to these aforementioned challenges, multiple new solutions have recently been proposed both in academia and in industry. Google's Explainable Artificial Intelligence (AI) Cloud [Goo20], for example, assists in the development of interpretable and explainable ML models and supports their deployment with increased confidence. Another example is the Descriptive machine Learning EXplanations (DALEX) [Dal20] package, which offers various functionalities that help users understand how complex models work. Some works propose to enable domain experts to collaborate with each other to tackle this problem together [CJH19, FBG19]. In this context, information visualization (InfoVis) techniques have been shown to be effective in making analysts more comfortable with ML solutions. Krause et al. [KPB14], for example, present a case study of domain experts using their tool to explore predictive models in electronic health records. Also, in visual analytics (VA), the first stages to partially address those challenges have already been reached, for instance by discussing how global [RSG16a] or local [MPG*14] interpretability can assist in the interpretation and explanation of ML [GBY*18, Wol19], and how to interactively combine visualizations with ML in order to better trust the underlying models [SSK*16].

We build our state-of-the-art report (STAR) upon the results of existing visualization research, which has emphasized the need for improved trust in areas, such as VA in general, dimensionality reduction (DR), and data mining. Sacha et al. [SSK*16] aimed to clarify the role of uncertainty awareness in VA and its impact on human trust. They suggested that the analyst needs to trust the outcomes in order to achieve progress in the field. Sedlmair et al. [SBIM12] found important gaps between the needs of DR users and the functionalities provided by available methods. Such limitations reduce the trust that users can put in visual inferences made using scatterplots built from DR techniques. Bertini and Lalanne [BL09] concluded, from a survey, that visualization can improve model interpretation and trust-building in ML. An interesting paper by Ribeiro et al. [RSG16b] shows that the interest on using visualization to handle issues of trust is also present in the ML field. The authors describe a method that explains the predictions of any classifier via textual or visual cues, providing a qualitative understanding of the relationship between the instance's components. Despite all the currently proposed solutions, many unanswered questions and challenges still remain, e.g., (1) If the analysts are not aware of the inherent uncertainties and trust issues that exist in an ML system, how to ensure that they do not form wrong assumptions? (2) Are there any guarantees that they will not be deceived by false (or unclear) results? (3) What problems of trustworthiness arise in each of the phases of a typical ML pipeline?

In this STAR, we present a general mapping of the currently available literature on using visualization to enhance trust in ML models. The mapping consists of details about which visualization techniques are used, what their reported effectiveness levels are, which domains and application areas they apply to, a conceptual understanding of what trust means in relation to ML models, and what important challenges are still open for research. Note that the

terms *trust* and *trustworthiness* are used interchangeably throughout the report. The main scientific contributions of this STAR are:

- an empirically informed definition of what trust in ML models means;
- a fine-grained categorization of trust against different facets of interactive ML, extracted from 200 papers from the past 12 years;
- an investigation of existing trends and correlations between categories based on temporal, topic, and correlation analyses;
- the deployment of an interactive online browser (see below) to assist researchers in exploring the literature of the area; and
- further recommendations for future research in visualization for increasing the trustworthiness of ML models.

To improve our categorization, identify exciting patterns, and promote data investigation by the readers of this report, we have deployed an interactive online survey browser available at

<https://trustmlvis.lnu.se>

We expect that our results will support new research possibilities for different groups of professionals:

- beginners/non-experts who want to get acquainted with the field quickly and gain trust in their ML models;
- domain experts/practitioners of any discipline who want to find the appropriate visualization techniques to enhance trust in ML models;
- model developers and ML experts who investigate techniques to boost their confidence and trust in ML algorithms and models; and
- early-stage and senior visualization researchers who intend to develop new tools and are in search of motivation and ideas from previous work.

The rest of this report is organized as follows (see Figure 1). In Section 2, we introduce background information that we used in order to comprehend the concept of trustworthiness of ML models. We also describe our adopted definition of the meaning of trust in ML models. In Section 3, we discuss existing visualization surveys that are relevant to our work. Afterwards, Section 4 provides details with regard to our methodology, i.e., the searched venues and the paper collection process. The overview in Section 5 includes initial statistical information. In Section 6, we present our categorization and describe the most representative examples. In Section 7, we report the results of a topic analysis performed on these papers to find new and interesting topics and trends derived from them, and further findings from data-driven analysis. Our interactive survey browser and research opportunities are discussed in Section 8. Finally, Section 9 concludes the STAR. Additionally, a set of supplementary materials (referred to as **S1** to **S8**) is also available, including the documents used to guide our categorization methodology, as well as the data that could not be part of this report due to space restrictions.

2. Background: Levels of Trustworthiness of Machine Learning Models

First, we present some earlier definitions of trust that are subsequently adapted to the context of our research. We also discuss

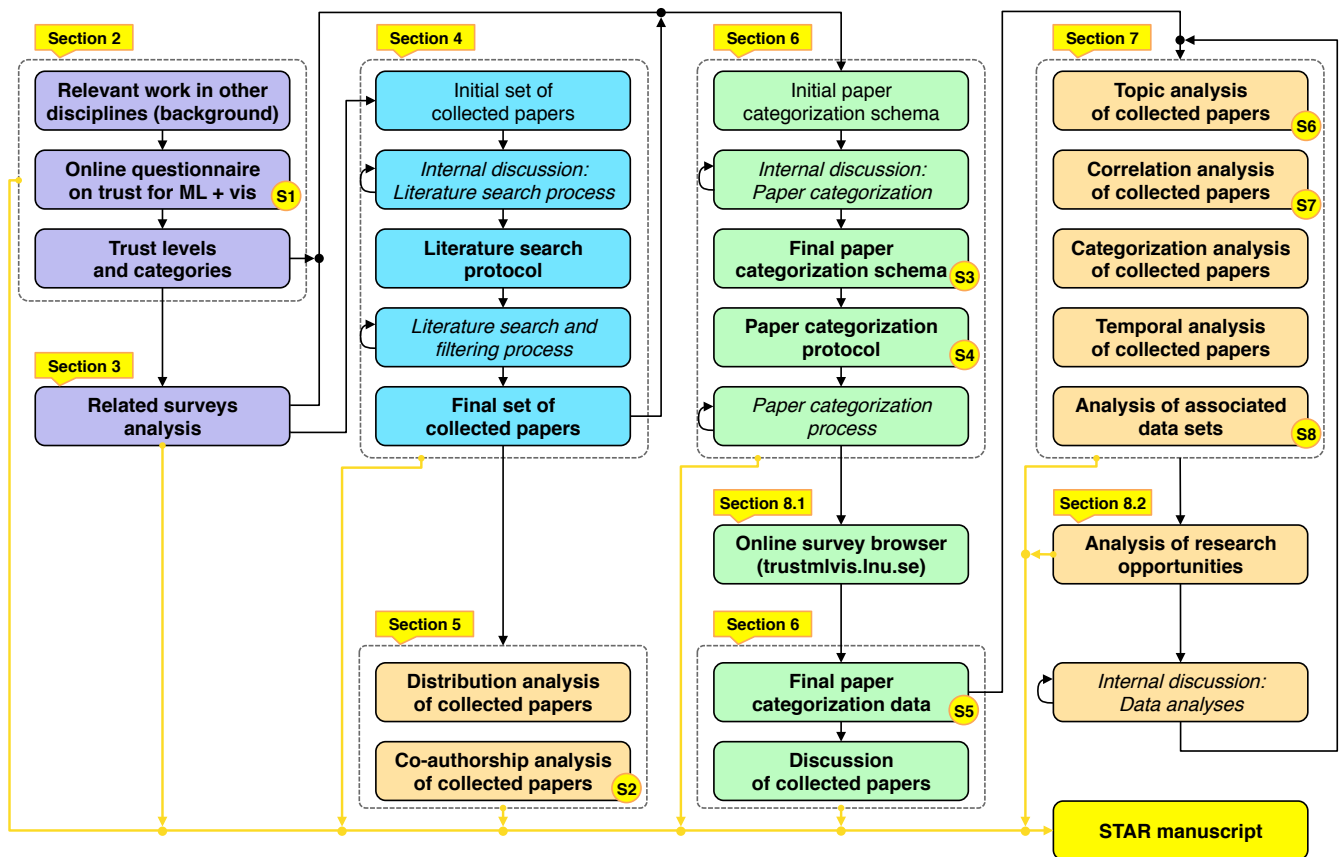


Figure 1: The overview of our STAR with regard to the methodology, main results, and corresponding sections of the manuscript. Color coding is used for grouping related activities and results (purple for the background information and key concepts, blue for the literature search, green for the paper categorization, orange for the data analyses, and yellow for the manuscript); italic font is used for intermediate activities; and bold font is used for the items discussed explicitly in this STAR. The marks S1–S8 refer to supplementary materials.

qualitative data gathered from an online questionnaire that we distributed among ML experts and practitioners. The goals of the questionnaire were to shape our categorization of trust issues in ML and to bring to light potential ideas on how visualization can support the improvement of trustworthiness in the ML process. Building upon these definitions and results, we group the identified factors of trust into five *trust levels* (TLs). These levels are a part of our overall methodology, discussed in Section 6.

Definitions of trust. The issues of definition and operationalization of *trust* have been discussed in multiple research disciplines, including psychology [EK09] and management [MDS95]. Such definitions typically focus on trust in the context of expectations and interactions between individuals and organizations. The existing work in human-computer interaction (HCI) extends this perspective. For example, Shneiderman [Shn00] provides guidelines for software development that should facilitate the establishment of trust between people and organizations. To ensure *trustworthiness* of software systems, he recommends the involvement of independent oversight structures [Shn20]. Fogg and Tseng [FT99] state that “trust indicates a positive belief about the perceived reliability of, dependability of, and confidence in a person, object, or

process”; in their work, trust is also related (and compared) to the concept of *credibility*. Rather than focusing on *interpersonal* trust, the existing work has also addressed *trust in automation* [HJBU13], which is more relevant to our research problem. Lee and See provide the following definition, widely used by the researchers in this context [LS04]: trust is “the attitude that an agent will help achieve an individual’s goals in a situation characterized by uncertainty and vulnerability”. This definition has been further extended by Hoff and Bashir [HB15], who propose a model of trust in automation with factors categorized into multiple *dimensions* and *layers*. Further adaptation of such multi-dimensional approach has been demonstrated, for example, by Yu et al. [YBC*18]. Lyons et al. [LHF*18] adopt a model consisting of a non-orthogonal set of factors in their analysis of trust factors for ML systems. **In this STAR**, we rely on the rather general definition of ML trust by Lee and See [LS04] and further expand it into a more detailed, multi-level model presented below. Additionally, we make use of the definitions and factors of trust described in the existing work within InfoVis and VA and incorporate them in our model. For example, Chuang et al. [CRMH12] define trust as “the actual and perceived accuracy of an analyst’s inferences”. Although important,

this definition touches only on a single perspective of trust: the one related to the users' expectations. The authors also mention that usually, during evaluations of the design choices performed on new visualization systems and tools, the modelling of choices and relationships *between* views is often omitted. This practice introduces limitations regarding the improvement of trust for the system as a whole, as opposed to the trustworthiness of each view individually. In the uncertainty typology detailed by MacEachren et al. [MRO*12], trust is decomposed into three high-level types: (i) *accuracy*, defined as correctness or freedom from mistakes, conformity to truth or to a standard or model; (ii) *precision*, defined as the exactness or degree of refinement with which a measurement is stated or an operation is performed; and (iii) *trustworthiness*, defined as source dependability or the confidence the user has in the information. The latter is a broad category that includes components such as completeness, consistency, lineage, currency/timing, credibility, subjectivity, and interrelatedness.

Online questionnaire. The next step of our work can be compared to domain problem characterization [Mun09]. In order to elicit the expectations and suggestions from ML practitioners with regard to our problem, we distributed an online questionnaire titled "How Would Visualization Help Enhancing Trust in Machine Learning Models?" (see supplementary material S1). We received answers from 27 participants, all with at least a Bachelor's degree, and most with a Master's (40.7%) or a Doctorate degree (51.9%). Almost all of them had their education in Computer Science or related fields. Some participants have only used ML in a few projects (around 33.3%), but most are either ML practitioners (22.2%) or developers/researchers in the field (44.4%). Their experiences with different types of ML algorithms/models are diverse, with rather balanced numbers between supervised (85.2%) and unsupervised (70.4%) learning. Within these two categories, classification (95.7%) and clustering (89.5%) are the most popular, respectively. The questionnaire itself begins with a description of a hypothetical scenario where a real-world data set was used (Pima Indians Diabetes, obtained from the UCI ML repository [DG17]). Each of the 15 questions presents a possible use of visualization related to trust in ML, and participants are asked to score them from 1 (strong disagreement) to 5 (strong agreement). Questions are also accompanied by short descriptions of some characteristics of the proposed scenario, in order to help participants in answering them.

According to the results, the bulk of the answers in most of the questions is concentrated around the scores of 4 and 5. This is evidence that the overall attitude of the participants towards visualization for enhancing trust in ML is largely positive. Factors such as visualizing details of the source of the data (Q1), data quality issues (Q3), performance comparison of different ML algorithms (Q4), hyper-parameter tuning (Q5), exploration of "what-if" scenarios (Q11), and investigation of fairness (Q12) obtained the majority of votes on score 5. Other factors which showed very positive—but less overwhelming—opinions were the visualization of details about the data collection process (Q2), data control and steering during the training process (Q6 and Q9), feature importance (Q7), visualizing the decisions of the model (Q8 and Q10), enabling collaboration (Q13), and the choice of tools for specific models (Q14). In these cases the majority of the scores were 4, but with some variance towards 3 and 5. The only question that deviated from

Table 1: Summary of the answers to the two open-ended questions on the participants' expectations and suggestions, which were provided at the end of the online questionnaire. The answers are sorted based on the number of occurrences # and then alphabetically.

Summary of open answers	#
Everything (i.e., all together)	5
Data importance (i.e., impact of data in the output)	4
Feature importance	4
Evaluation	3
Impact of hyper-parameters	3
Pre-processing	3
Evolution of metrics during training	2
Investigation of outliers	2
Model accuracy and results in general	2
Raw data	2
Class separation	1
Decision boundaries	1
Learned features	1
Uncertainty	1

this trend was the last one (Q15), where we proposed that a single well-designed performance metric would be enough to judge the quality of an ML model, and no further actions (such as visualization) would be necessary. In this case, most of the scores were concentrated on either 1 or 2, showing clear disagreement.

The questionnaire ends with two open-ended questions, where participants were free to give their ideas and opinions on which steps of the ML process (or properties of the models and the data) they would like to visualize to increase the trust in the ML models they use. Many participants indicated their desire to visualize the ML process as much as possible, in all phases where it might apply (5 answers). Additionally, out of all the specific concepts and ideas that emerged, the most popular were the visualization of feature importance (4 answers), the impact of different characteristics of the data instances (4 answers), investigation of hyper-parameters (3 answers), visualizing the pre-processing steps (3 answers), and the evaluation of the model (3 answers). Table 1 summarizes all these answers along with the number of occurrences. These answers were mostly aligned with our prior hypotheses, but also enabled us to gain new insights on what was missing from our categorization of trust factors (see below). For instance, the *source reliability* category was influenced by one participant who described her/his work with Parkinson's disease data and the reliability problems involved with it: "For instance, I have been working with clinical studies with Parkinson's disease patients wearing sensors in their wrists. For us researchers, it was difficult to see how the data was collected e.g. patients could do a certain daily activity (e.g. cutting grass) but in our model we accounted that as tremor." Another important point that was brought up is that visualization-based steering of the ML training process might push the user to "fish" for desired results and invalidate the statistical significance of the model.

Trust levels (TLs) and categories. In this STAR, we cover the subject of *enhancing trust in ML models* with the use of *visualizations*. As such, we do not cover solutions proposed to address those questions solely at the algorithmic level, even if they are considered with growing interest by ML researchers (as exemplified by the two plenary invited talks [How18, Spi18] on the subject given in 2018 at NeurIPS, one of the major ML venues). Based on the

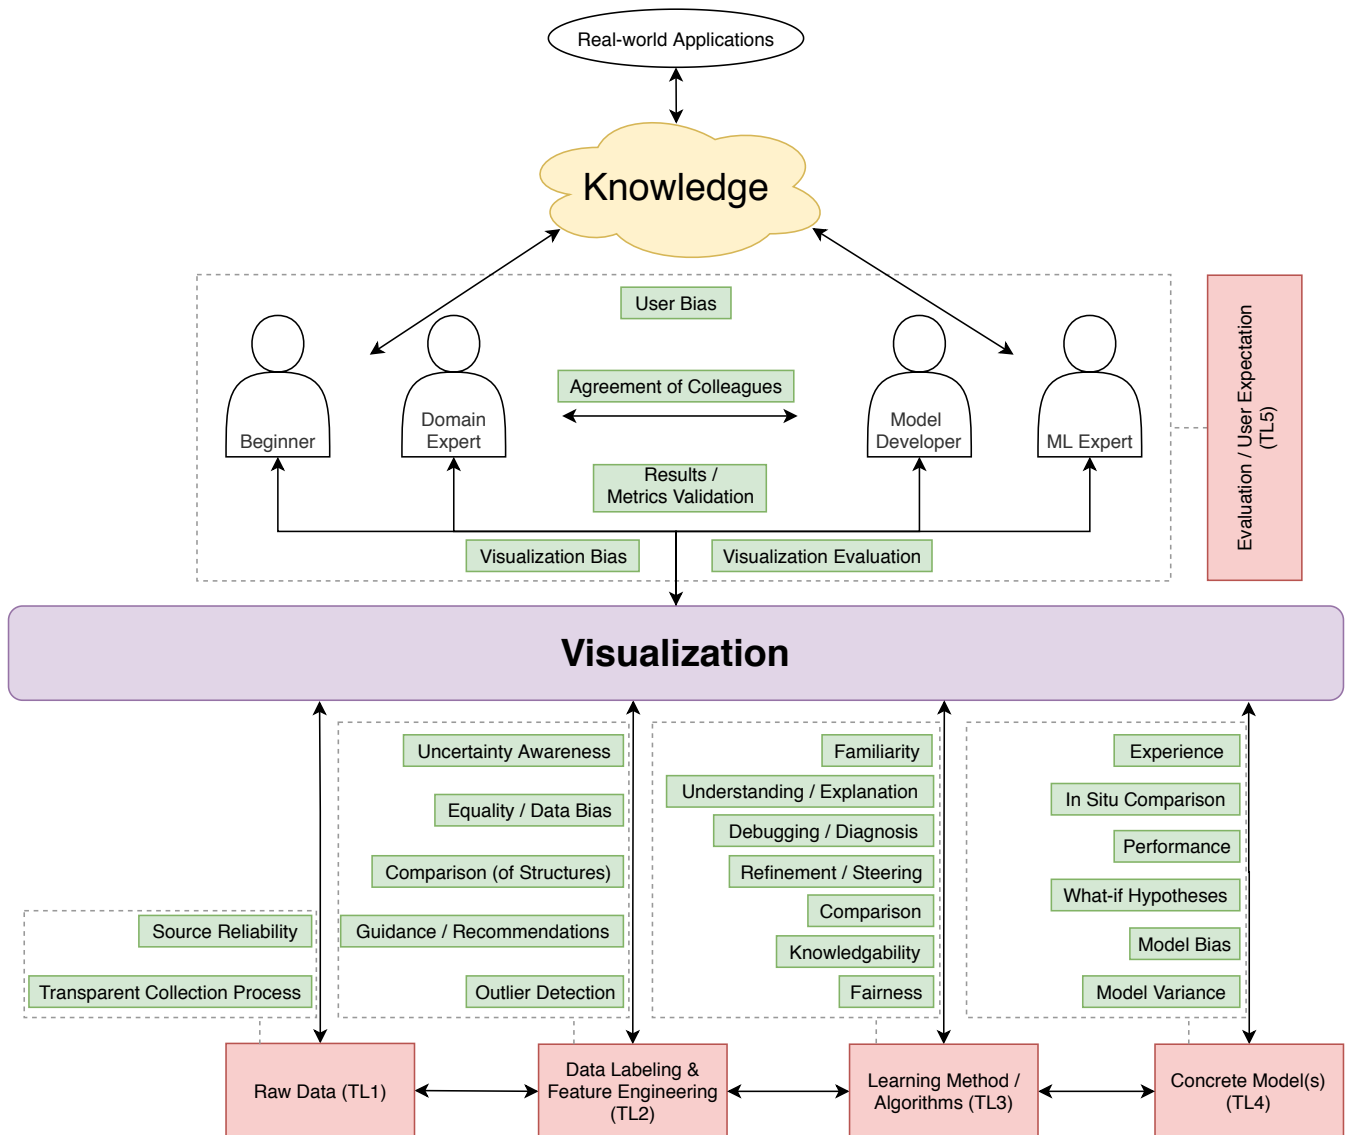


Figure 2: A typical ML pipeline (depicted in red), assisted by visualization (in purple). Issues of trust permeate the complete shown pipeline, and we locate and categorize these issues in several **trust levels (TLs)**. The various categories proposed in this work are represented in green. The yellow “cloud” represents the knowledge created by the different target groups while they pursue their goals by using visualizations to explore the pipeline, the data and/or the ML models. Finally, at the very top, we encode the real-world applications with an ellipsoid.

existing work discussing the issues of trust, the suggestions from ML experts (see above), and internal discussions, we consider that the problem of enhancing trust in ML models has a multi-level nature. It can be divided into five **TLs** related to trustworthiness of the following: *the raw data* (→TL1), *the processed data* (→TL2), *the learning method (i.e., the algorithms)* (→TL3), *the concrete model(s) for a particular task* (→TL4), and *the evaluation and the subjective users’ expectations* (→TL5). These levels of trust are aligned with the usual data analysis processes of a typical ML pipeline, such as (1) collecting the raw data; (2) allowing the user to label, pre-process, and query/filter the data; (3) interpreting, exploring, and explaining algorithms in a transparent fashion; (4) re-

fining and steering concrete model(s); and (5) evaluating the results collaboratively. With the term *algorithm*, we define an ML method (e.g., logistic regression or random forest); in contrast to a *model* which is the result of an algorithm and is trained with specific parameters.

We use the term *level* to refer to the increasingly abstract nature of concepts as well as to emphasize the sequential aspect of the ML pipeline. Indeed, the lack of trustworthiness in each stage of the pipeline cumulatively introduces instability in the predictions of a model. Thus, trust issues (i.e., categories) that are relevant to two or more of our TLs are assigned to the lowest TL possible. This is similar to concerns about issues cascading from earlier to later

levels within the nested model for visualization design and validation [Mun09], for instance. Figure 2 displays the connection between a typical ML pipeline (slightly adapted from the work of Sun et al. [SLT17]) and the visualization techniques that enhance trust in ML models in various phases. In the bottom layer of Figure 2, we depict (in red) the ML pipeline comprising the distinct areas where users are able to interact with (and choose from) a large pool of alternatives or combinations of options. The layer above depicts the visualization (in purple). The upper layer consists of the different target groups that we address, the generation of knowledge, and the usability of this knowledge in solving problems stemming from real-life applications. Finally, the multiple categories of trust associated with each of these levels are presented in green in Figure 2 and discussed in detail below.

- **Raw data (TL1).** The lowest trust level gathers categories attached to the data collection itself. They belong to the complex task of preparing the data for further analysis, commonly referred to as *data wrangling* [KHP*11].

Arguably, *source reliability* is the very first category that should be visualized in a system. It should detect and handle the cases that do not meet the quality expectations or that show unusual behavior. For instance, detecting that some labels are unreliable could guide the user in selecting ML algorithms that are resistant to label noise [FV14]. However, perceiving source reliability is not an easy task, as it involves visualization questions, such as “how to visualize the data source involved in data collection?”, but also the very statistical questions of measuring reliability.

As a proxy for this measure, one can visualize information, for instance, “was a particular university involved in data collection, was a domain expert such as a doctor present during the health data collection, and were the sensors reliable and error-free?” Hence, source reliability is strongly related to ensuring a *transparent collection process*, the second category of this level. This includes visualizing the data collection process, what systems were used to collect the data, and how, why, and how objectively that was done.

Issues about reliability of the data and of the collection process can jeopardize, from the very start, the ML process and diminish the TLs set by users. If those issues remain undetected, they can spoil the later phases, according to the classic “garbage in, garbage out” principle. For instance in the case of unreliable labels [FV14], reported error rates are also unreliable. This is becoming more relevant with the growing attention given to adversarial machine learning, an ML research field which focuses on adversarial inputs designed to break ML algorithms or models [GMP18, LL10].

- **Data labeling & feature engineering (TL2).** The next group of categories has its focus beyond the raw data and into feature engineering and labeling of the data. This is also partially related to *data wrangling*. Trust issues at this level focus on data that are overall considered to be reliable and clean. Trust can then be enhanced by addressing subgroup or instance problems. With *uncertainty awareness* and visualizations supporting it, the data instances that do not fit can be filtered out, and any borderline cases are highlighted to be explored by the users via visual representations.

The category *equality/data bias* is related to the *fairness* category discussed below. It concerns the possible sources of subgroup-specific bias in the decision of an ML model. For instance, if a subgroup of the population has characteristics that are significantly different from the ones of the population as a whole, then the decisions for members of this subgroup could be unfair compared to the decisions for members of other subgroups. Visualization methods can be used to explore interesting subgroups and to pinpoint potential issues.

Comparison (of structures) [KCK17] implies the usage of visualization techniques in order to compare different structures in the data. As an example, experts in the biology domain would like to compare different structures visually, and furthermore, improve these representations with various encodings such as color.

Guidance/recommendations [CGM19] is a good continuation of the previous concept: trust can be improved by using visualization tools that (1) recommend new labels in the unlabeled data scenarios, for example, in semi-supervised learning and (2) guide the user to manage the data by adding, removing, and/or merging data features and instances.

Finally, for this level of trust, *outlier detection*, i.e., searching and investigating extreme values that deviate from other observations of a data set, can be alleviated by visualization systems (this is a major issue in ML [CBK09]). Detecting and manipulating in a meaningful way an observation that diverges from an overall pattern on a sample is a useful way to positively influence the results and boost overall trust in the process. Notice that this category focuses on particular instances, while the *source reliability* category described previously, considers data globally.

- **Learning method/algorithms (TL3).** This group of categories concerns the ML algorithms themselves, as the third step of the ML pipeline. Each category corresponds to a particular way of enabling a better control, in broad sense, over ML algorithms.

Familiarity is how visualization can support users in order to help them getting familiar with a certain learning method. There is a possibility that users are biased towards using an ML algorithm they know instead of the others that might actually be more appropriate. Improvement of familiarity by using visualization could help to both limit this type of bias and to enhance the users’ trust in algorithms they do not know well.

Interpretability and explainability are among the most common and widespread categories—being found in most of the papers that we identified. We further subdivide both into the following categories:

- understand the reasons behind ML models’ behavior and why they deviate from each other (*understanding/explanation*);
- diagnose causes of unsuccessful convergence or failure to reach a satisfactory performance during the training phase (*debugging/diagnosis*);
- guide experts (and novices) to boost the performance, transparency, and quality of ML models (*refinement/steering*); and
- compare different algorithms (*comparison*).

It should be noted that the issue of interpretability and explainability has been receiving a growing attention in the ML commu-

nity. Algorithms are modified in order to produce models that are easier to interpret. However, those models are frequently claimed to be more interpretable based on general rules of thumb, such as “rule-based systems are easier to understand than purely numerical approaches” or “models using fewer features than others are easier to understand”. Only the most recent papers tend to include user-based studies [AGW18, CSWJ18]. Unfortunately, they only explore quite simple visualization techniques such as static scatterplots.

Knowledgeability translates to the question: if users are not aware of an ML algorithm, then how are they supposed to use it? Possible solutions to provide assistance to users in such situations include visualizations designed to compare different models or to provide details about each algorithm. However, the lack of visualization literacy limits the possibilities for exploration of an ML algorithm and effects negatively all the categories of this phase [BRBF14, BBG19]. Model-agnostic (more general) visualization techniques that consider multiple algorithms can also support this challenge.

Last but not least, the category of *fairness* covers the analysis of subgroup-specific effects in ML prediction, e.g., whether predictions are equally accurate for various subgroups; for instance, females versus males, or if there are discrepancies that give a group an advantage or a disadvantage compared to other groups. This topic has recently received a lot of attention in the ML community. It has been shown in particular that the most natural fairness and performance criteria are generally incompatible [KMR17]. Thus, ML algorithms must make some compromises between those criteria which justify the strong need for visually monitoring/analyzing such trade-offs.

- **Concrete model(s) (TL4).** This final step of the ML pipeline consists of turning its inputs, mainly a set of ML learning methods/algorithms, into a *concrete model or a combination of models* [SKKC19]. Trust issues related to this step concern mostly performance related aspects, both in a static interpretation but also in a dynamic/interactive way.

Experience is a primary crucial factor, since promoting personalized visualizations based on the experiences of a user alter and might determine the selection. As an example, an expert in ML, a novice user, and a specific domain expert have different needs, and “what are their experiences and how can the visualization adapt to that?” is an important question.

In situ comparison can be described as comparing different snapshots and/or internal structures of the same concrete model in order to enhance trust.

Performance is another very common way to monitor the results of a model visually. Performance can objectively compare a model with another. However, this is usually insufficient for a complete understanding of the trade-off between different models.

What-if hypotheses appear when users search for impacts based on their interactions. A potential question is: “What is the consequence if we change one parameter and keep the rest stable for a specific model, or select some points to explore further?”

Model bias and *model variance* are well-known concepts originating from statistics with regard to the bias-variance trade-off. The *bias* is a systematic error created by the wrong hypotheses in

a model. High bias can cause a model to avoid seeing the relevant associations between features and target outputs, thus underfitting. The *variance* is a manifestation of the model’s sensitivity or the lack thereof to the data, more precisely to the training set. It could also be the result of parameterizations or perturbations. High variance can result in a model which bears inside random noise in the training data, rather than the intended outputs, hence overfitting.

- **Evaluation/user expectation (TL5).** The last group of categories is subsequent to visualization tools and techniques that oversee the ML pipeline, leading to knowledge generation in the overall workflow. Evaluation of models and meeting user expectations [CRMH12] is a key component for people to trust or not ML model(s) for a task.

Agreement of colleagues is supported by visualizations with provenance [OAB*17, RESC16] and collaborative visualizations which facilitate, for instance, ten experts from diverse domains to agree that a model performed well. This purpose could be served by provenance features and specific glyphs or snapshots, along with web-based online tools and platforms. When using visualizations, the choices of the visual metaphor and the visual variable (e.g., color instead of size) are important but can supplement the process negatively with *visualization bias*. This kind of bias was described, for example, by Lespinats and Aupetit [LA11]. However, this issue is being addressed by multiple ongoing research efforts in various subfields of visualization which are outside of the scope of this survey [MHSW19, XPGF19]. Thus, we have not included this perspective in our categorization.

A measure against visualization bias that we consider instead is the *visualization evaluation* [Mun09] that many authors of visualization papers perform. Quantitative or qualitative methods are used in the InfoVis and other communities to evaluate new visualization techniques. Both count as visualization evaluations, even receiving feedback from ML experts and/or domain experts before, during, or after the development of a visualization system.

Moreover, *results/metrics validation* is the most common method utilized by developers of visualization tools to indicate if a model can be trusted and has reached user expectations. However, we believe that it is not sufficient on its own.

Finally, *user bias* is a rarely addressed category which tries to understand the cognitive biases of users who have the power to steer an automated process. Questions such as where, when, and why a user has to interact with a model are still an open challenge. A paper from Nalcaci et al. [NGB*19], for example, works with distinction bias and confirmation bias in visualization systems that are related to user bias when viewing visualizations. Also, a recent survey from Dimara et al. [DFP*20] tries to connect the possibly-biased judgment and decision making of humans with specific visualization tasks.

3. Related Surveys

The challenge of enhancing trust in ML models has not yet received the same level of attention in systematic surveys as other topics, for example, the understanding and interpretation of DR or deep learn-

ing (DL). To the best of our knowledge, this is the only survey that deals with InfoVis and VA techniques for enhancing the trustworthiness of ML models. In order to confirm that, we carefully examined the online browser of the survey of surveys (SoS) paper from McNabb and Laramée [ML17], which contains 86 survey papers from main InfoVis venues. We have also investigated 18 additional survey papers in our own recent SoS paper [CMJK20]. Our analysis indicated that many of these surveys are about interpretable ML models, especially regarding current popular subjects such as *interpretable and interactive ML (IML)*, *predictive VA (PVA)*, *DL*, and *clustering and DR*. None of these papers, however, has an explicit focus on categorizing and/or analyzing techniques related specifically to the subject of trust in ML models. Related issues, such as accuracy, quality, errors, stress levels, or uncertainty in ML models, are touched upon by some of them, but in our work these issues are discussed in more detail. Particularly, uncertainty in the data and the visualization itself is a part of our TLs in the *uncertainty awareness* and the *visualization bias* categories. One of the main differences in our work is the focus on the transformation from uncertainty to trust, which should happen progressively in all phases of an ML pipeline. Some previous works offer brief literature reviews and propose frameworks for human-centered IML with visualization [SSZ*16, SSZ*17], the problem of integrating ML into VA [ERT*17], trust in VA [SSK*16], or comparison of DR techniques from an ML point of view [VDMPVdH09]. Although interesting, those papers fall outside the scope of the *trust in ML models* subject. One of the motivations for this STAR came from our analysis of the future work sections of these surveys—10 out of the 18 surveys highlight the subject of enhancing trust in the context of ML models, making this challenge one of the most emergent and essential to solve. This body of work also forms the basis for our methodological part of the literature research, presented in Section 4.

3.1. Interpretable and Interactive Machine Learning

The work concerning the interpretability of ML models in the visualization community started to emerge around 15 years ago. This opportunity was captured by Liu et al. [LWLZ17] who conducted a survey that summarizes several ML visualization tools focusing on three categories (*understanding*, *diagnosis*, and *refinement*). This is different compared to our perspective and goal to categorize only those papers that tackle the problem of enhancing trust in ML models. The recent publication by Du et al. [DLH20] groups techniques for interpretable ML into *intrinsic* and *post-hoc*, which can be additionally divided into *global* and *local* interpretability. The authors also suggest that these two types of interpretability bring several advantages, for example, that users trust an algorithm and a model's prediction. However, they do not analyze in details the different aspects of enhancing trust in ML models as we performed in this STAR. Overall, these surveys (and the categories from Liu et al. [LWLZ17], together with *comparison*) target the interpretability and explainability at the level of ML algorithms, which are themes under the umbrella of VIS4ML (visualization for ML) and comprise only a small subset of our proposed categorization.

Moreover, the topic of IML aided by visualizations has been discussed in many papers recently, as it was summarized in the surveys

by Amershi et al [ACKK14] and Dudley and Kristensson [DK18]. The former focused on the *role of humans in IML* and *how much users should interfere and interact with ML*. They also suggested at which stages this interaction could happen and categorized their papers accordingly. Steering, refining, and adjusting the model with domain knowledge are not trivial tasks and can introduce cumulative biases into the process. Due to this, in this STAR our analysis focuses on the biases that a user might introduce into a typical ML pipeline. Furthermore, visualizations may introduce different biases to the entire process, as discussed in the previous Section 2. In such situations, the visualization design should be directed towards conveying, or occasionally removing, any of these biases initially and not simply making it easier for users to interact with ML models.

3.2. Predictive Visual Analytics

Lu et al. [LGH*17] adopted the pipeline of PVA, which consists of four basic blocks: (i) *data pre-processing*, (ii) *feature selection and generation*, (iii) *model training*, and (iv) *model selection and validation*. These are complemented by two additional blocks that enable interaction with the pipeline: (v) *visualization* and (vi) *adjustment loop*. The authors also outline several examples of quantitative comparisons of techniques and methods before and after the use of PVA. However, no analysis has been performed about trust issues that are incrementally added in each step of the pipeline. Another survey written by Lu et al. [LCM*17] follows a similar approach by classifying papers utilizing the same PVA pipeline, but with two new classes: (a) *prediction* and (b) *interaction*. For instance, regression, classification, clustering, and others are the primary subcategories of the prediction task; and explore, encode, connect, filter, and others, are subcategories of interaction. This work inspired us to introduce the *interaction technique* subcategory of our basic category, called *visualization*. One unique addition, though, is the *verbalize* category, which describes how visualization and use of words can assist each other by making the visual representation more understandable to users and vice versa. Concluding, none of these survey papers so far provide future opportunities touching the subject of how visualization can boost ML models' trustworthiness.

3.3. Deep Learning

Grün et al. [GRNT16] briefly explain how the papers they collected are separated to their taxonomy for feature visualization methods. The authors defined three discrete categories as follows: (i) *input modification methods*, (ii) *deconvolutional methods*, and (iii) *input reconstruction methods*. Undoubtedly, learned features of convolutional neural networks (CNNs) are a first step to provide trust to users for the models. But still, this step belongs to the interpretability and explainability of a specific algorithm, i.e., very specialized and targeted to CNNs. In our work, we cover not only CNNs but every ML model with a focus on the data, learning algorithms, concrete models, users, and thus not only on the model. The two main contributions of Seifert et al. [SAB*17] are the analysis of insights that can be retrieved from deep neural network (DNN) models with the use of visualizations and the discussion about the visualization techniques that are appropriate for each type of insight. In their paper, they surveyed visualization papers and distributed them into

five categories: (1) *the visualization goals*, (2) *the visualization methods*, (3) *the computer vision tasks*, (4) *the network architecture types*, and (5) *the data sets that are used*. This paper is the only one that contains analyses for the data sets used in each visualization tool, which worked as a motivation for us to include a data set analysis in our survey. However, their main contributions do not touch the problem of trustworthiness, but more the correlation of visualizations and patterns extraction (or insights gaining) for DNNs. A summarization of the field of interpreting DL models was performed by Samek et al. [SWM18], putting into the center the increasing awareness of how interpretable and explainable ML models should be in real life. The main goal of their survey is to foster awareness of how useful it is to have interpretable and explainable ML models. General interpretability and explainability play a role in increasing trustworthiness, but not a major one. The different stages of the ML pipeline should be taken into account as from early stages, bias and deviance can occur and grow when processing through the pipeline. Zhang and Zhou [ZZ18] study their papers starting from the visualization of CNN representations between network layers, over the diagnosis of CNN representations, and finally examining issues of disentanglement of “the mixture of patterns” of CNNs. They neither provide a distinct methodology of categorization for their survey, nor insights on the problem of trust as opposed to our survey.

Another batch of papers on DL assembles into Garcia et al.’s [GTdS*18] survey in which visualization tools addressing the interpretability of models and explainability of features are described. The authors focus on various types of neural networks (NNs), such as CNNs and recurrent neural networks (RNNs), by incorporating a mathematical viewing angle for explanations. They emphasize the value of VA for the better understanding of NNs and classify their papers into three categories: (a) *network architecture understanding*, (b) *visualization to support training analysis*, and (c) *feature understanding*. In a similar sense, (i) *model understanding*, (ii) *debugging*, and (iii) *refinement/steering* are three directions that Choo and Liu [CL18] consider. Model understanding aims to communicate the rationale behind model predictions and spreads light to the internal operations of DL models. In cases when the DL models underperform or are unable to converge, debugging is applied to resolve such issues. Finally, model refinement/steering refers to methods that enable the interactive involvement of usually experienced experts who build and improve DL models. Compared to our survey, only half of the learning methods are considered. Thus, their reader support is limited when it comes to show how their algorithms actually work on several occasions. Yu and Shi [YS18] examined visualization tools that support the user to accomplish four high-level goals: (1) *teaching concepts*, (2) *assessment of the architecture*, (3) *debugging and improving models*, and (4) *visual exploration of CNNs, RNNs, and other models*. They describe four different groups of people in their paper: (a) *beginners*, (b) *practitioners*, (c) *developers*, and (d) *experts*, distributed accordingly to the four aforementioned classes. These groups are also considered in our work. Nonetheless, teaching concepts and assessing the architectures of DNNs are particular concepts that do not enhance trust explicitly. This is why we focus on multiple other categories, such as models’ *trade-off of bias and variance* or *in situ comparisons of structures of the model*, in general and not

exclusively for DL models. Hohman et al. [HKPC19] surveyed VA tools that explore DL models by investigating papers into six categories answering the aspects of “who”, “why”, “what”, “when”, “where”, and “how” of the collected papers. Their main focus is on interpretability, explainability, and debugging models. The authors conclude that just a few tools visualize the training process, but solely consider the ML results. Our *ML processing phase* category is motivated by this gap in the literature, i.e., we investigate this challenge in our paper to gain new insights about the correlation of trust and visualization in *pre-processing*, *in-processing*, and *post-processing* of the overall ML processing phases. Finally, as many explainable DL visualization tools incorporate clustering and DR techniques to visualize DL internals, the results of these methods should be validated on how trustworthy they are.

3.4. Clustering and Dimensionality Reduction

Sacha et al. [SZS*17] propose, in their survey, a detailed categorization with seven guiding scenarios for interactive DR: (i) *data selection and emphasis*, (ii) *annotation and labeling*, (iii) *data manipulation*, (iv) *feature selection and emphasis*, (v) *DR parameter tuning*, (vi) *defining constraints*, and (vii) *DR type selection*. During the *annotation and labeling* phase, for example, hierarchical clustering could assist in defining constraints which are then usable by DR algorithms. Nonato and Aupetit [NA19] separate the visualization tools for DR according to the categories *linear and nonlinear*, *single- versus multi-level*, *steerability*, *stability*, and *others*. Due to the complexity of our own categorization and our unique goals, we chose to use only their first category (*linear versus nonlinear*), as is common in previous work [VDMPVdH09]. Nonato and Aupetit also describe different quality metrics that can be used to ensure trust in the results of DR. However, as the results of our online questionnaire suggested (cf. Section 2), comparing those quality metrics alone is probably not sufficient. To conclude, the main goal of these two surveys is not related to ML in general, and the latter one only discusses trust in terms of aggregated quality metrics. This is a very restricted approach when compared to our concept of trust, which should be ensured at various levels, such as *data*, *learning method*, *concrete model(s)*, *visualizations themselves*, and *covering users expectations*.

4. Methodology of the Literature Search

In the following, we present the methodology used to identify and systematically structure the papers of our STAR. Our work is inspired by the same methodology guidelines from Lu et al. [LGH*17], Garcia et al. [GTdS*18], and Sacha et al. [SZS*17] presented in Section 3. In an initial pilot phase (cf. [Sny19]), we extracted appropriate keywords from ten relevant papers [VSK*15, WJCC16], including those that deal with the problems of interpretable/explainable ML (which are closely related to trust in ML). The keywords were divided into two lists with the goal to cover both *trust* and *ML*. For *trust*, the used keywords were, in alphabetical order: “accuracy”, “assess”, “bias”, “black box”, “confidence”, “diagnose”, “distort”, “error”, “explain”, “explore”, “feedback”, “guide”, “interact”, “noise”, “quality”, “robustness”, “stress”, “trust”, “uncertainty”, “validate”, “verify”, and their

derivatives. For *ML*, the searched keywords were: “artificial intelligence”, “classification”, “clustering”, “deep learning”, “dimensionality reduction”, “machine learning”, “neural network”, “projections”, and all the types of *ML* (e.g., “supervised learning”).

The keywords from the two lists were combined into pairs, such that each keyword from the first list was paired with each keyword from the second. These paired keywords were used for seeking papers relevant to the focus of this survey in different venues (cf. Section 4.1). A validation process was used in order to scan for new papers and admit questionable cases, as described in Section 4.2. Papers that were borderline cases and eventually excluded are discussed in Section 4.3.

4.1. Search and Repeatability

To gather our collection of papers, we manually searched for papers published in the last 12 years (from January 2008 until January 2020). We started our search from InfoVis journals, conferences, and workshops, and later extended it to well-known *ML* venues (the complete list can be found at the end of this subsection). Moreover, when seeking papers in *ML*-related venues (e.g., the International Conference on Machine Learning, ICML), we included two additional keywords: “visual” and “visualization”.

Within the visualization domain, we checked the following resources for publications:

Journals: IEEE TVCG, Computers & Graphics (C&G), Computer Graphics Forum (CGF), IEEE Computer Graphics & Applications (CG&A), Information Visualization (IV), DiStill, and Visual Informatics (VisInf).

Conferences: IEEE Visual Analytics in Science and Technology (VAST), IEEE Visualization Conference (VIS) short papers track, Eurographics Visualization (EuroVis), IEEE Pacific Visualization (PacificVis), ACM Conference on Human Factors in Computing Systems (CHI), and ACM Intelligent User Interfaces (IUI).

Workshops: Visualization for AI Explainability (VISxAI), EuroVis Workshop on Trustworthy Visualization (TrustVis), International EuroVis Workshop on Visual Analytics (EuroVA), Machine Learning Methods in Visualisation for Big Data (MLVis), Visualization for Predictive Analytics (VPA), Visual Analytics for Deep Learning (VADL), IEEE Large Scale Data Analysis and Visualization (LDAV), and Visualization in Data Science (VDS).

Within the *ML* domain, we checked the following venues:

Conferences: ICML, Knowledge Discovery and Data Mining (KDD), and European Symposium on Artificial Neural Networks, Computational Intelligence, and Machine Learning (ESANN).

Workshops: ICML Workshop on Visualization for Deep Learning (DL), ICML Workshop on Human Interpretability in *ML* (WHI), KDD Workshop on Interactive Data Exploration & Analytics (IDEA), NIPS Workshop on Interpreting, Explaining and Visualizing Deep Learning.

The search was performed in online libraries, such as IEEE Xplore, ACM Digital Library, and Eurographics Digital Library. As an example of the number of results we got, both IEEE Transactions

on Visualization and Computer Graphics (TVCG) and IEEE Visual Analytics in Science and Technology (VAST) together resulted in around 750 publications. Due to the use of a couple of broad keyword combinations in order to cover our main subject effectively, some of the papers collected were not very relevant. They were sorted out in the next phase of our methodology.

4.2. Validation

For the sake of completeness, we quickly browsed through each individual paper’s related work section and tried to identify more relevant papers (a process known as *snowballing* [Woh14]). With this procedure, we found more papers belonging to other venues, such as the Neurocomputing Journal, IEEE Transactions on Big Data, ACM Transactions on Intelligent Systems and Technology (ACM TIST), the European Conference on Computer Vision (ECCV), Computational Visual Media (CVM), and the Workshop on Human-In-the-Loop Data Analytics (HILDA), co-located with the ACM SIGMOD/PODS conference. In more detail, this validation phase was performed in four steps:

1. we removed unrelated papers by reading the titles, abstracts, and investigating the visualizations;
2. we split the papers into two categories: *approved* and *uncertain*;
3. uncertain papers were reviewed by at least two authors, and if the reviewers agreed, they were moved to the approved papers;
4. for the remaining papers (i.e., where the two reviewers disagreed), a third reviewer stepped in and decided if the paper should be moved to the approved category or discarded permanently.

The calculated amount of disagreement, i.e., the number of conflicts in the 70 uncertain cases, was less than 20% (approximately 1 out of 5 papers). This process led to 200 papers that made it into the survey.

4.3. Borderline Cases

We have restricted our search to papers with visualization techniques that explicitly focus on supporting *trust* in *ML* models, and not on related perspectives (e.g., assisting the exploration and labeling process of input data with visual means). Therefore, papers such as those by Bernard et al. [BHZ*18, BZSA18], Gang et al. [GRM10], and Kucher et al. [KPSK17], although undoubtedly interesting, are out of the scope of our survey, since their research contributions are exclusively based on labeling data. Other partially-related papers [AASB19, AW12, BHGK14, FBT*10, SBTK08, ZSCC18] are also not included because they focus on using clustering solely to explore the data, without addressing inherent problems of the method. For similar reasons, the paper by Wenskovitch et al. [WCR*18], that tries to connect and aggregate benefits from clustering and DR methods, was excluded. Moreover, papers on high-dimensional data clustering or exploratory data analysis are not included (e.g., Behrisch et al. [BKSS14], Lehmann et al. [LKZ*15], Nam et al. [NHM*07], and Wu et al. [WCH*15]). Finally, there are related works that provide important contributions to the visualization community, but do not study trust explicitly, and thus were not included:




Figure 3: Histogram of the set of collected techniques/tools (200 in total) with regard to the publication year. (*) Please note that the data for 2020 is incomplete since the data collection for this survey was completed in January 2020. (†) For 2007, we did not perform a complete search; the single publication was found within the related work section of another already-included paper.

improving the computational performance of algorithms (e.g., t-SNE) [PHL*16, PLvdM*17], frameworks and conceptual designs for closing the loop [SKBG*18], investigating cognitive biases with respect to users [CCR*19], and enabling collaboration with the use of annotations [CBY10].

5. General Overview of the Relations Between the Papers

This section begins with a meta-analysis of the spatiotemporal aspects of our collection of papers. The analysis shows, on the one hand, that there is an increasing trend in trust-related subjects; on the other hand, it also highlights the struggles of collaborations between visualization researchers and ML experts. Additionally, we generated a co-authorship network to observe the connections of the authors from all the papers. By exploring the network and its missing links, we hope to bring researchers closer to form new collaborations towards research in the trustworthiness of ML models.

Time and venues. Our collection of papers comprises 200 entries from a broad range of journals, conferences, and workshops. The analysis of the temporal distribution (see Figure 3) shows a stable growth in interest in the topic since 2009, with a sharp increase in 2018 and 2019 (and promising numbers also for 2020). The numbers for the publication venues identified can be seen in Table 2. Visualization researchers seem to be very interested in working with solutions to this problem and try to extend their work in ML venues with the creation of new workshops. There is a large number of workshops on the topic, co-located with ML venues, which indicates that researchers are interested in reaching out of their respective areas in order to collaborate. However, the small number of publications outside of visualization venues could possibly show a struggle of visualization researchers to find and collaborate with ML experts. It might also indicate that ML experts are not fully aware of the possibilities that the visualization field provides.

Co-authorship analysis. We analyzed the co-authorship network of the authors of our collection of papers using

Table 2: Number of visualization techniques # with regard to the respective publication venues in visualization (left and middle columns) and other disciplines (mostly ML venues; right column). Journals are marked with ‘J’ and workshops with ‘W’. The remaining venues are conferences.

Visualization venue	#	Visualization venue	#	Other disciplines venue	#
IEEE TVCG J	56	EuroVis	3	WHI W (ICML)	3
CGF J	33	VADL W (VIS)	3	ACM TIST J	2
IEEE VAST	24	MLVis W (EuroVis)	3	ESANN	2
ACM IUI	10	C&G J	2	HILDA W (ACM SIGMOD)	2
EuroVA W (EuroVis)	9	Distill J	2	Neurocomputing J	1
ACM CHI	8	IEEE VIS	2	IEEE Tran. on Big Data J	1
IV J	6	VISxAI W (VIS)	2	KDD	1
IEEE CG&A J	5	VisInf J	1	ECCV	1
IEEE PacificVis	5	CVM J	1	DL W (ICML)	1
VPA W (VIS)	4	LDAV W (VIS)	1	IDEA W (KDD)	1
VDS W (VIS)	4	TrustVis W (EuroVis)	1		

Gephi [BHO9], as presented in Figure 4. The goal was to identify a potential lack of collaboration within the visualization and ML communities. Enhancing collaboration between specific groups may lead to improvements in the subject of boosting trust in ML models with visualizations. The more connections an author has, the bigger is the size of the resulting node, i.e., the in-degree values of the graph nodes are represented by node size in the drawing. We colored the top eight clusters with the highest overall in-degree for all the nodes of each cluster. Finally, we filtered the node labels (authors first names and surnames) by setting a limit to the in-degree value in order to reduce clutter. By looking at the resulting co-authorship network (see Figure 4 and S2), we can observe a huge cluster in violet ①. In this cluster, Huamin Qu, Remco Chang, Daniel A. Keim, Cagatay Turkay, and Nan Cao seem to be the most prominent authors, with many connections. If we consider different subclusters in this massive cluster, Nan Cao is the bridge between some of the subclusters. Another cluster on the left (with light green color ②) is related to the big industries (such as Google and Microsoft) with Fernanda B. Viégas, Martin Wattenberg, and Steven M. Drucker as the most eye-catching names. Interestingly, this industry cluster is very well separated from the remaining academic clusters. Potentially, the connection of this industry cluster with the remaining clusters could have an impact on the research output produced by the visualization community. There are many smaller clusters of collaborating people, for example, the cluster with David S. Ebert and Wei Chen ③, Klaus Mueller ④, Han-Wei Shen ⑤, Alexandru C. Telea ⑥, Valerio Pascucci ⑦, and others (e.g., ⑧) obviously serving as main coordinators.

6. In-Depth Categorization of Trust Against Facets of Interactive Machine Learning

In this section, we discuss the process and results of our categorization efforts. We introduce a multifaceted categorization system with the aim to provide insights to the reader about various aspects of the data and ML algorithms used in the underlying literature. The main sources of input for the categorization were the previous work from the surveys discussed in Section 3, the iterative process of selecting papers (and excluding the borderline cases) described in Section 4, and the feedback received from the online questionnaire (Section 2). The top two levels of the proposed hierarchy of categories can be seen below, with 8 overarching aspects (6.1 to

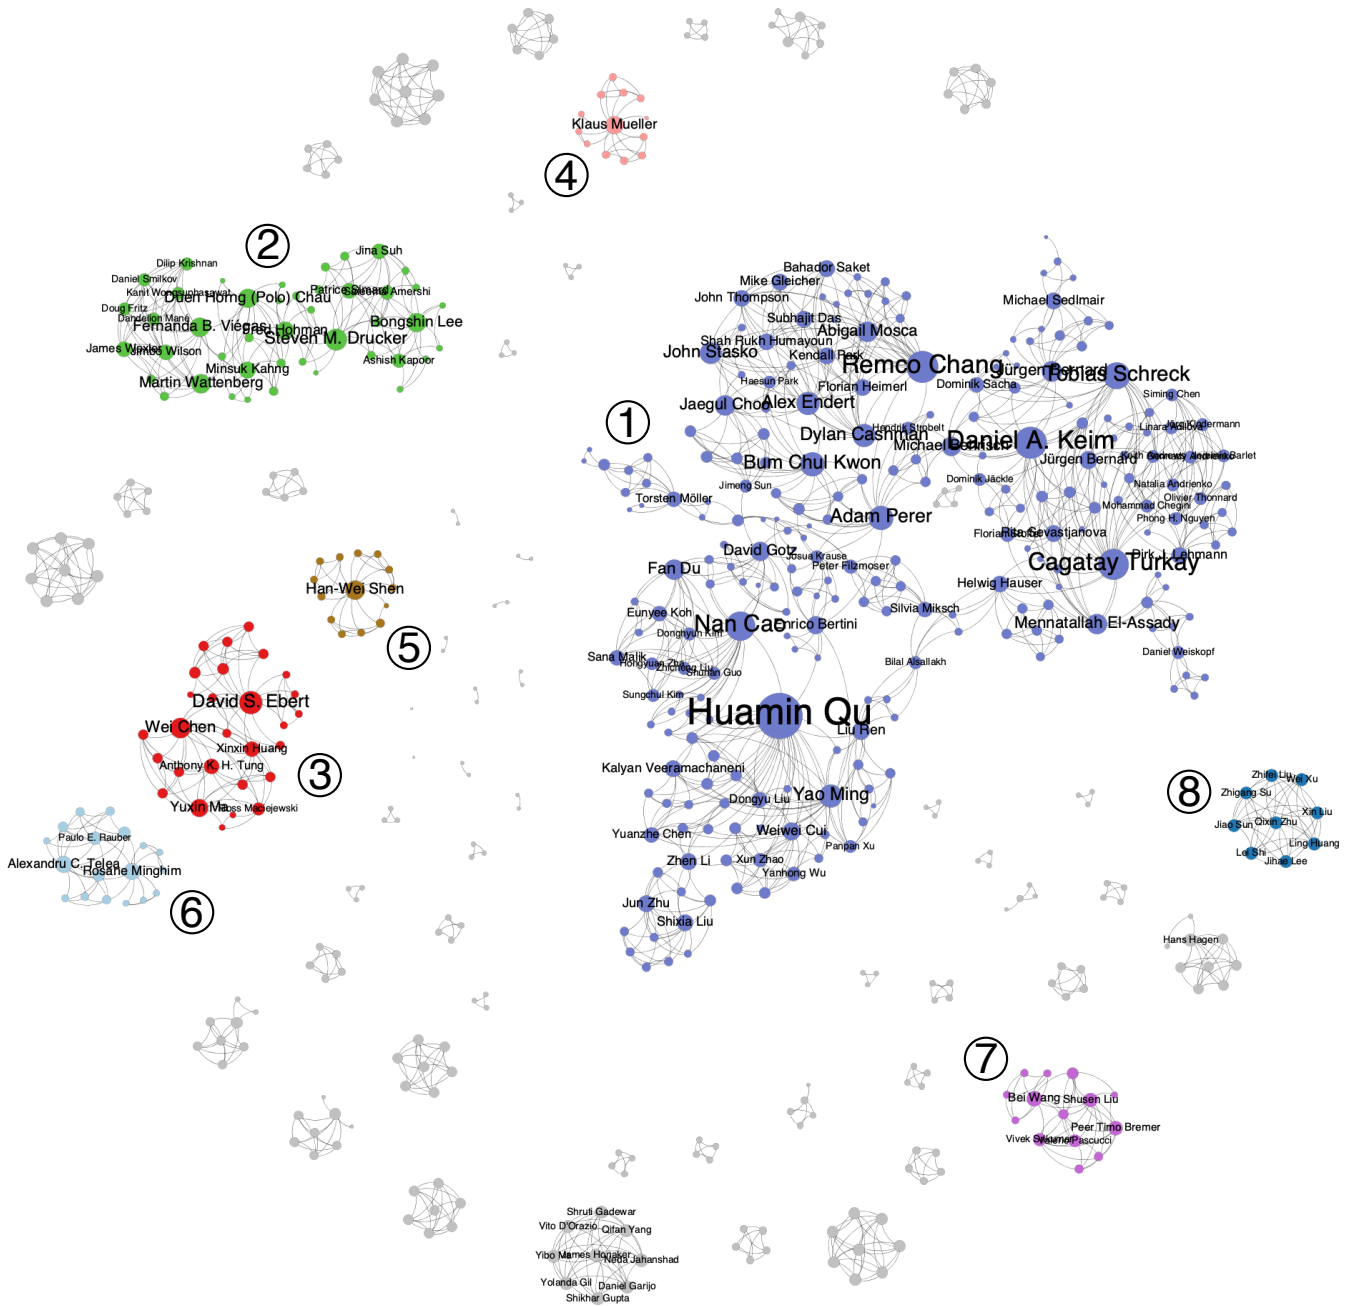


Figure 4: Co-authorship network visualization with the eight largest connected components (①–⑧) highlighted in different colors. The node size represents the in-degree centrality of each author. The labels are filtered based on the in-degree value in order to reduce clutter.

6.8), partitioned into 18 category groups (6.1.1 to 6.6.3, plus TL1 to TL5), resulting in a total of 119 individual categories.

- 6.1. Data
 - 6.1.1. Domain (10 categories)
 - 6.1.2. Target Variable (5 categories)
- 6.2. Machine Learning
 - 6.2.1. ML Methods (16 categories)

- 6.2.2. ML Types (10 categories)
- 6.3. ML Processing Phase (3 categories)
- 6.4. Treatment Method (2 categories)
- 6.5. Visualization
 - 6.5.1. Dimensionality (2 categories)
 - 6.5.2. Visual Aspects (2 categories)
 - 6.5.3. Visual Granularity (2 categories)
 - 6.5.4. Visual Representation (19 categories)

- 6.5.5. Interaction Technique (9 categories)
- 6.5.6. Visual Variable (6 categories)
- 6.6. Evaluation
 - 6.6.1. User Study (2 categories)
 - 6.6.2. User Feedback (2 categories)
 - 6.6.3. Not Evaluated (1 category)
- 6.7. Trust Levels (TL) 1–5
 - TL1: Raw Data (2 categories)
 - TL2: Processed Data (5 categories)
 - TL3: Learning Method (7 categories)
 - TL4: Concrete Model (6 categories)
 - TL5: Evaluation/User Expectation (4 categories)
- 6.8. Target Group (4 categories)

The complete overview of all categories is shown in Table 3 (also in S3 as a mind map). The aspect and category group names are preceded by the subsection numbers where they are introduced and discussed. This is to avoid confusion and to reduce the cognitive load of the reader.

Designing the categorization. Compared to previous surveys, we created new categories to better cover the 200 papers we included in our STAR. In the following list, we present the basic purpose for each aspect, along with the core similarities and differences when compared to the related surveys from Section 3.

- **6.1. Data** tries to create a link between the input data/application and the enhancement of trust in ML models. The first category group we identified in this aspect is the data *domain*. We took the inspiration from our previous publication [KPK18], but the categories are significantly different to fit the current subject. In the case of the *target variable*, we conceived the idea of separating the independent variable of each data set.
- **6.2. Machine Learning** is an inherent component in boosting the trustworthiness of ML models. We used several sources for defining parts of the *ML methods* category group. Different neural networks methods, such as CNNs, RNNs, DCNs, and DNNs, were contained in other works [SAB*17, YS18]. Also, as mentioned in Section 3, linear and non-linear DR is an existing categorization from [NA19], but ensemble learning and the remaining DL categories are new to our STAR. *ML types*, such as classification, regression, and clustering, can be seen in the work [LGH*17], but we improved this short categorization by using the complete supervised, unsupervised, semi-supervised, and reinforcement learning distinction.
- **6.3. ML Processing Phase** connects the ML and the visualization aspects and shows when VA techniques are deployed to improve the trustworthiness of the ML models. *During* and *after training* categories can be found in the work of Hohman et al. [HKPC19], which in our case are adjusted to the newly introduced *pre-processing*, *in-processing*, and *post-processing* phases.
- **6.4. Treatment Method** deals with differences between *model-agnostic* or *model-specific* approaches. Observing such distinctions might indicate where the community should later focus on to better boost the trust in ML models. *Model-agnostic vs. model-specific* methods used in VA systems are first described in our work, although Dudley and Kristensson [DK18] hinted about *model agnosticism*.

- **6.5. Visualization** is another inherent component of how increasing trust in ML models can be achieved. Visualization details, such as *dimensionality*, can also be found in the work of Kucher et al. [KK15]. However, we added *visual aspects and granularity*. *Visual representation* was also inspired by Kucher et al. [KK15] and many of the other related surveys. The *verbalize* category is a novel addition to pre-existing work, which is part of the 6.5.5. *Interaction Technique* group described in the work of Lu et al. [LCM*17]. Finally, for this aspect, the work of Kucher et al. [KPK18] covers all the visual variables used by us except for *opacity*.
- **6.6. Evaluation** of visualization can reduce the visualization bias, thus boosting even further the application of VA systems for ML. We are among the first who included this new aspect to highlight the importance of evaluations in visualization systems, tools, and techniques.
- **6.7. Trust Levels (TL) 1–5**, introduced in Section 2, is the most novel category group. Only TL3, which contains the interpretability/explainability group of categories, is described in previous works [LWLZ17, CL18]. Despite that, *comparison* is a fresh addition to this group.
- **6.8. Target Group** is equally important to the problem of enhancement of the trustworthiness in ML models as the input (i.e., the data) and the actual visualization. This aspect is inspired mainly by Yu and Shi's paper [YS18].

Overall, this extensive categorization aims to completely unveil the relationship between trust and the remaining categories, as can be seen later in Figure 6 and Section 7.

Filling in the categorization. To ensure consistency during the process of assigning the 200 papers to our categories in a first cycle, we created a code of practice (see S4) as a base structure. This base structure provides guidance to evaluate the individual papers in the same way without misalignment between the authors of this STAR. We also cleared and double-checked the resulting data for any issues that could come up with the annotated data in a second cycle. In particular, we looked for the following issues: (1) outliers, (2) typos, (3) discrepancies, and (4) inconsistencies between different evaluators by inspecting and removing any obscure and misclassified data cases. The fact that a large subset of the papers (75%) were classified by the same author also maximizes the consistency of their final categorization. Each classified paper belongs to zero, one or more categories for every aspect depending on the information the paper contains. Due to the page limits and readability concerns, we cannot discuss all 200 papers in this section. Instead, we only focus on the most prominent and (in our opinion) most important ones. All 200 papers are referenced in Table 4 and are part of the bibliography. The complete survey data set, including the individual categorization for each paper, is provided in our online survey browser (see Fig. 8) and in S5.

6.1. Data

Many visualization techniques have been tested with specific data sets coming from different domains. However, just a few of them specifically work for one type of data set, for instance, the systems proposed by Bremm et al. [BvLBS11] and Wang et al. [WGZ*19]. In this subsection, we present the most frequent data domains we

Table 3: The categorization used in our survey. The total number of corresponding visualization techniques per category is shown in every row, along with heatmap-style icons.

6.1.1. Domain	198
Biology	28
Business	19
Computer Vision	59
Computers	6
Health	30
Humanities	41
Nutrition	8
Simulation	8
Social / Socioeconomic	22
Other	93

6.1.2. Target Variable	199
Binary (categorical)	39
Multi-class (categorical)	128
Multi-label (categorical)	9
Continuous (regression problems)	24
Other	38

6.2.1. ML Methods	198
Convolutional Neural Network (CNN)	25
Deep Convolutional Network (DCN)	8
Deep Feed Forward (DFF)	10
Deep Neural Network (DNN)	19
Deep Q-Network (DQN)	9
Generative Adversarial Network (GAN)	10
Long Short-Term Memory (LSTM)	13
Recurrent Neural Network (RNN)	18
Variational Auto-Encoder (VAE)	14
Other (DL methods)	21
Linear (DR)	57
Non-linear (DR)	51
Bagging (ensemble learning)	27
Boosting (ensemble learning)	11
Stacking (ensemble learning)	6
Other (generic)	97

6.2.2. ML Types	197
Classification (supervised)	111
Regression (supervised)	20
Other (supervised)	7
Association (unsupervised)	5
Clustering (unsupervised)	41
Dimensionality Reduction (unsupervised)	66
Classification (semi-supervised)	13
Clustering (semi-supervised)	6
Classification (reinforcement)	1
Control (reinforcement)	3

Legend: 0 papers 100 papers 200 papers

6.3. ML Processing Phase	198
Pre-processing / Input	36
In-processing / Model	45
Post-processing / Output	162

6.4. Treatment Method	196
Model-agnostic / Black Box	144
Model-specific / White Box	70

6.5.1. Dimensionality	199
2D	196
3D	5

6.5.2. Visual Aspects	199
Computed	195
Mapped	109

6.5.3. Visual Granularity	200
Aggregated Information	183
Instance-based / Individual	146

6.5.4. Visual Representation	199
Bar Charts	82
Box Plots	11
Matrix	50
Glyphs / Icons / Thumbnails	63
Grid-based Approaches	19
Heatmaps	46
Histograms	56
Icicle Plots	6
Line Charts	56
Node-link Diagrams	47
Parallel Coordinates Plots (PCPs)	32
Pixel-based Approaches	8
Radial Layouts	22
Scatterplot Matrices (SPLOMs)	18
Scatterplot / Projections	115
Similarity Layouts	27
Tables / Lists	86
Treemaps	5
Other	59

6.5.5. Interaction Technique	185
Select	163
Explore / Browse	169
Reconfigure	74
Encode	112
Filter / Query	113
Abstract / Elaborate	177
Connect	128
Guide / Sheperd	48
Verbalize	9

6.5.6. Visual Variable	196
Color	195
Opacity	83
Position / Orientation	58
Shape	37
Size	68
Texture	17

6.6. Evaluation	200
Standard	38
Comparative	12
Before / During Development	36
After Development	47
Not Evaluated	102

6.7. Trust Levels (TL) 1-5	200
Source Reliability	11
Transparent Collection Process	6
Uncertainty Awareness	29
Equality / Data Bias	16
Comparison (of Structures)	86
Guidance / Recommendations	46
Outlier Detection	68
Familiarity	3
Understanding / Explanation	95
Debugging / Diagnosis	54
Refinement / Steering	69
Comparison	61
Knowledgeability	10
Fairness	6
Experience	7
In Situ Comparison	54
Performance	108
What-if Hypotheses	40
Model Bias	19
Model Variance	16
Agreement of Colleagues	9
Visualization Evaluation	87
Metrics Validation / Results	130
User Bias	7

6.8. Target Group	196
Beginners	41
Practitioners / Domain Experts	162
Developers	36
ML Experts	73

spotted and the nature of the target variable that should be predicted by ML classifiers.

6.1.1. Domain

DeepVID, by Wang et al. [WGZ*19], is an example that focuses only on images (i.e., the overall field of **computer vision**). It shows

how ML models for image classification can be interpreted and debugged with the use of simpler models (e.g., a linear model) in DNNs. Since DNNs usually work well with image data, exploring and diagnosing the training process of a DNN is an initial step towards boosting trust in them. With regard to **humanities** data, Sherkat et al. [SNMM18] propose a system that empowers users

incorporating their feedback to define several diverse algorithms for clustering. The supported interactions enable users to adjust (or even create) new key terms, which are then used to supervise the algorithms and cluster the different documents. Involving humanities experts in user studies to evaluate the effectiveness of VA systems can further increase trust. An example of operating with **health** data is INFUSE [KPB14], which helps analysts to select features and retrieve extra information based on a selection from a collection of algorithms, cross-validations folds, and different ML models. The visual representations assist domain experts (i.e., doctors) in manipulating their medical records more precisely and improve the accuracy of the results. In this case, the visualization tool seems to be generalizable to other domains. In medicine, receiving recommendations during the data processing phase is necessary to ensure that no further biases are introduced in the input phase of an ML model. This guidance is better achievable by feature exploration and feature selection with the use of visualization. Bremm et al. [BVLBS11] focus on **biological** data sets. In their paper, the authors utilize scatterplot- and grid-based visualizations to facilitate selection and later comparison of data descriptors for unlabeled biology-related data. Except for the comparison of data and structures, any uncertainties stemming from the data should be highlighted for the biologists to focus on them. As a consequence, this extraction of patterns through visualization can increase their trust in ML. A number of papers focus on various other data domains, such as the works of Gleicher [Gle13], Sips et al. [SNLH09], and Tatu et al. [TMF*12].

6.1.2. Target Variable

In ML, the target (or response) variable is the characteristic known during the learning phase that has to be predicted for new data by the learned model. In classification problems, it can take a *binary* value for two-class problems, a single label for *multi-class* problems, or even a set of labels for *multi-label* problems. In regression problems, it is generally a *continuous* variable. On the one hand, Krause et al. [KDS*17] propose a workflow to help practitioners to examine, diagnose, and explain the decisions made by a binary classifier. In their approach, instance-level explanations are obtained based on local feature significance measures that explain single instances. With these findings as a basis, they develop visualizations that lead the users to important areas of investigation. Extensions of this approach could evaluate the reliability of the incoming data by comparing different areas and timeframes of the data. On the other hand, a multi-class data set has been used in ActiVis [KAKC18]. This visualization tool integrates coordinated multiple views, such as a graph that provides an overview of the model architecture, and a neuron activation view for exploring DNN models with user-defined subsets of instances. To wrap up the decisions made based on the provided views, agreement of colleagues could further enhance the trust in these complex DNNs. Many approaches exist for regression problems, such as those described in the publications of Fernstad et al. [FSJ13] and Hohman et al. [HSD19]. For instance, Piringer et al. [PBK10] describe a validation framework for regression models that enable users to compare models and analyze regions with poor predictive performances. The optimization of the so-called *trade-off of bias and variance* is also crucial for regression problems.

Papers categorized as *others* on the target variable group concern ML settings in which no target variable is available. These are mostly related to DR and clustering problems. The method designed by Zhou et al. [ZLH*16], for example, combines both aspects to enable users to design new dimensions from data projections of subspaces, with the goal of maintaining important cluster information. The newly adapted dimensions are included in the analysis together with the original ones, to help users in forming target-oriented subspaces that explain—as much as possible—cluster structures.

6.2. Machine Learning

This subsection covers various ML methods that were divided into three main classes: *DL*, *DR*, and *ensemble learning*. We also discuss different ML types that we considered in our categorization: *supervised learning*, *unsupervised learning*, *semi-supervised learning*, and *reinforcement learning*.

6.2.1. Machine Learning Methods

In the area of *DL*, we observed two categories that are in the focus of most DL-related papers: *CNNs* and *RNNs*. CNNComparator [ZHP*17] addresses the challenges of comparing CNNs and enables users to freeze the tool for different epochs of a trained CNN model (by using so-called snapshots). An epoch is completed when a data set is processed one time forward and backward through an NN. Thus, CNNComparator provides insights into the architectural design, and as a consequence, it enables better training of CNN models. For RNNs, RNNbow [CPM*18] visualizes the gradient flow while backpropagation occurs in the training of RNNs. By visualizing the gradient, this tool offers insights into how exactly the network is learning. Both papers explicitly enhance trust in different DL models by either comparing CNNs or explaining RNNs to the users via visualization. In the *DR* subclass, linear techniques surpass non-linear ones in volume (the former was found in 57 papers vs. 51 for the latter). One example here is the iPCA tool [JZF*09]. It augments the principal component analysis (PCA) algorithm with an interactive visualization system that supports the investigation of relationships between the data and the computed eigenspace. Overall, the tool employs views for exploring the data, the projections, the PCA's eigenvectors, and the correlations between them. Eventually, prominent uncertainties become aware to users by examining all these relations. In another example, AxiSketcher [KKW*17] enables users to impose their domain knowledge on a visualization by allowing interaction using a direct-manipulation technique over a t-SNE projection (non-linear DR technique). Users can sketch lines over specific data points, and the system composes new axes that represent a non-linear and weighted mixture of multidimensional attributes. Thus, the comparison of clusters enables users to identify problematic cases (in terms of trust) in a projection. For *ensemble learning*, *bagging* is the most common category. iForest [ZWLC19] is a visualization tool that accommodates users with an aggregated view showing and summarizing the decision paths in random forests, which finally reveals the working mechanism of the ML model. Visualizing and understanding the decision paths of random forest algorithms, as well as how their performance was reached, serves as a foundation for assessing the trust in bagging ensemble learning. Other,

more general examples can be found in the works by Schneider et al. [SJS*18] and Sehgal et al. [SRG*18].

6.2.2. Machine Learning Types

According to our analysis, *supervised learning* and *classification* problems are extensively addressed by the visualization community. For instance, a visualization system that works with choosing the best classifiers is EnsembleMatrix [TLKT09]. It allows users to directly interact with the visualizations in order to explore and build combinations of models. Comparison of ML models and validation metrics are key factors in increasing trust in them. HyperMoVal [PBK10], already discussed in Section 6.1.2, focuses on *regression* problems and provides several functionalities: comparing the ground truth against predicted results, analyzing areas with a poor fit, evaluating the physical plausibility of models, and comparing various classifiers. When users address regression problems, the comparison of alternative ML models and steering each of them can also improve the trustworthiness of the models. In *unsupervised learning*, a *clustering* example is the visualization technique developed by Turkey et al. [TPRH11], which visualizes the structural quality of several temporal clusters at a certain point in time or over time. DimStiller [IMI*10] is a system (belonging to the DR subclass) that assists the user in converting the input dimensions in a number of analytical steps into data tables that can be transformed into each other with the help of so-called operators. Users can manipulate those operators for parameter tuning and for guidance to discover patterns in the local neighborhood of the data space. Both DR and clustering visualization tools often utilize comparison of structures and emphasize patterns observable in projections. Some rare cases are related to *semi-supervised learning*, such as MacInnes et al. [MSW10] and Bernard et al. [BZL*18]. *Reinforcement learning* is covered by the work of Saldanha et al. [SPBA19], for instance.

6.3. Machine Learning Processing Phase

VASSL [KKZE20] is a system that works with the *pre-processing/input phase* and enhances the performance and scalability of the manual labeling process by providing multiple coordinated views and utilizing DR, sentiment analysis, and topic modeling. The system allows users to select and further investigate batches of accounts, which supports the discovery of spambot cases that may not be detected when checked independently. For the *in-processing/model phase*, Liu et al. [LSL*17] designed a tool that helps to better understand, diagnose, and steer deep CNNs. They represent a deep CNN as a directed acyclic graph, and based on this representation, a hybrid visualization has been developed to disclose multiple aspects of each neuron and the intercommunications between them. The largest category with regard to the number of available visualizations is *post-processing/output* for visualizing the final results, such as MultiClusterTree [VL09]. In their tool, the authors propose a 2D radial layout that supports an inherent understanding of the distribution arrangement of a multidimensional multivariate data set. Unique clusters can be explored interactively by using parallel coordinates when being selected in a cluster tree representation. The overall cluster distribution can be explored, and better understanding of the relations between clusters and the initial attributes is supported as well. As expected, the

input phase is highly related to TL2, the in-processing phase to understanding and steering categories of TL3, and the final phase to metrics validation (TL5). Finally, Gil et al. [GHG*19] and Sacha et al. with VIS4ML [SKKC19] are two workflow papers that provide an overview of all these phases.

6.4. Treatment Method

Model-agnostic techniques are twice as common as *model-specific* techniques. With the former, we mean—in most of the cases—visualization methods that treat ML models as *black boxes*. The latter is usually connected to techniques specifically developed to open these black boxes, and thus make the ML models to be regarded as *white boxes*. An example of a model-agnostic visualization tool is ATMSeer [WMJ*19] with which users are able to steer the search space of AutoML and explain the results. A multi-granular visualization empowers users to observe the AutoML process, examine the explored ML models, and refine the search space in realtime. In the white box case, the visualization tool EasySVM [MCM*17] facilitates users in tuning parameters, managing the training data, and extracting rules as a component of the support vector machine (SVM) training process. The goal of model-specific techniques is to explain the inner workings of a particular ML model. However, some tools combine both specific models and model-agnostic algorithms, such as Chae et al. [CGR*17], Roesch and Günther [RG19], Pezzotti et al. [PHV*18], and others [CWGW19, KKB19, MCMT14].

6.5. Visualization

Various approaches, types, and properties of visualization are used in our 200 surveyed papers, often in combinations. The knowledge of the most common techniques and approaches can guide early-stage researchers to choose the most important of them or senior researchers to discover potential gaps in the literature. The selection of the best visualization approaches, types, and properties for a given situation can effectively reduce potential visualization bias. Successfully addressing questions such as “where, when, and why should I use a 2D bar chart to present aggregated information instead of another visual representation?”, for instance, can boost trust in ML models. Carefully thinking about which data should be visualized is similarly important. This section of our report describes all these aspects and introduces the corresponding category groups.

6.5.1. Dimensionality

With regard to dimensionality of the visual display, almost all visualizations (196) are 2D, such as [dBD*12, JJ09, MXC*20, SDMT16]. An exception is the interactive visualization technique by Coimbra et al. [CMN*16] that adapts and improves biplots to show the data attributes in the projected *three-dimensional (3D)* space. They use interactive bar chart legends to present variables that are visible from a given angle and also support users to decide on the optimal position to examine a desired set of attributes.

6.5.2. Visual Aspects

The information to be visualized can either be directly mapped from the data values themselves or be computed (algorithmically

derived). ModelTracker [ACD*15] extracts information contained in conventional summary statistics and charts while letting users examine errors and diagnose ML models. Hence, it contains *computed* instead of *mapped* instances. Arendt et al. [ASW*19] visualize the classifier's feedback after each iteration with their IML interface. To address scalability issues of the visualization, this interface communicates with the user by a small set of system-proposed instances for each class.

6.5.3. Visual Granularity

Görtler et al. [GSS*20] represent *aggregated information* in their visualizations. They use a technique that performs DR on data that is subject to uncertainty by using a generalization of standard PCA. Their technique helps to discover high-dimensional characteristics of probability distributions and also supports sensitivity analysis on the uncertainty in the data. Zeiler and Fergus [ZF14] introduce a visualization technique that contributes to insight generation for the general operation of the classifier in an *instance-based* manner, i.e., for individual data cases. Nevertheless, most visualization systems and techniques involve both the exploration of aggregated information and individual cases, e.g., presented by Choo et al. [CLKP10] and the visualization tool BaobabView [vv11].

6.5.4. Visual Representation

Liu et al. [LXL*18] combine multiple coordinated views to provide a thorough overview of a tree boosting model and enable the effective debugging of a failed training process. One of their views utilizes *bar charts* in order to rank the most valuable features that affect the model's performance. Ji et al. [JSR*19] propose visual exploration of a neural document embedding with the goal to gain insights into the underlying embedding space and encourage this utilization in standard infrared (IR) spectroscopy applications. In their paper, they use a *scatterplot* visualization, i.e., a projection. LSAView [CDS09] is a system for interactive, latent semantic analysis (LSA) models. Multiple views, linked matrix-graph views, and data views in the form of *lists* are used to choose parameters and see the effects of them. Other papers apply different visual representations, some rare cases are *waterfall charts*, *violin charts*, *Voronoi diagrams*, and *bipartite graphs* [Aup07, HSD19, KKZE20, LA11, LGG*18, WGSY19, WGYS18, ZTR16].

6.5.5. Interaction Technique

Gehrmann et al. [GSK*20] argue that both the visual interface and model architecture of DL systems need to consider the interaction design. They propose a collaborative semantic inference for the constructive cooperation between humans and algorithms. Semantic interactions permit a user both to understand and regulate parts of a model's reasoning process. All these interactions enable the *selection* of particular sentences and then further *exploration* of the content with suggestions stemming from the system side. *Abstract/elaborate* is another interaction technique found in, e.g., Borland et al. [BWZ*20] and can be interpreted as different granularities that the visualization allows users to explore the data. Sevastjanova et al. [SBE*18] argue that a combination of visualization and *verbalization* methods is advantageous for generating wide and versatile insights into the structure and decision-making processes

of ML models. For more details about the remaining interaction techniques (e.g., [CD18a, CBB*19, CLKP10, PSMD14, PLHL19, ZWC*18]), we refer to the survey of Lu et al. [LGH*17].

6.5.6. Visual Variable

Ahmed et al. [AYMW11] use a qualitative *color* scheme in order to encode cluster groupings in all views of their visualization tool for steering mixed-dimensional KD-KMeans clustering. Color is used in almost every paper we examined [CSG*18, EASKC18, KS12, ML14, PSF17, XCH*16]. DeepCompare [MMD*19] uses *opacity* and *size*, which are the two second-most occurring *visual variables*. Their tool visualizes the results of DL models, provides insights into the model behavior and the assessment of trade-offs between two such models. In more detail, the activation value of an NN is encoded as *size*, while *opacity* is used to remove the highlighting when specific cases are selected.

6.6. Evaluation

In this subsection, we explore how visualizations are evaluated in our community and how many of them had been actually evaluated. Surprisingly, around half of the visualizations were never evaluated. The evaluation of visualizations is a fundamental component to validate the usability of visualization tools and systems.

6.6.1. User Study

RuleMatrix [MQB19] is one of the approaches that follows a *standard* procedure for performing an evaluation in the InfoVis community. That is, various participants had to solve a series of tasks by using the tool during which the accuracy and timing was monitored to gain insight into the usability of the proposed solution. The paper presents an interactive visualization technique to assist novice users of ML to understand, examine, and verify the performance of predictive models. FairSight [AL20] is another tool designed to accomplish different concepts of fairness in ranking decisions. To achieve that, FairSight distinguishes the required actions (understanding, computing, and others) that can possibly lead to fairer decision making. It was *compared* against the What-If Tool [WPB*20] and found to perform better and result in more benefits than the latter approach.

6.6.2. User Feedback

Cashman et al. [CHH*19] worked with exploratory model analysis, which is defined as the process of finding and picking relevant models that can be used to create predictions on a data source. *During development*, they improved their tool and received user *feedback*. Hazarika et al. [HLW*20] used networks as surrogate models for visual analysis, and *after the development* of their system and techniques, a domain expert gave them feedback in order to further improve the VA system at the end of the development process. Ultimately, from the further analysis of the statistics, we conclude that in five cases both domain and ML experts used visualization tools and evaluated them. In 32 cases, e.g., [CS14, KCK*19, LLL*19, MvW11, XYC*18], only domain experts were asked; and in 19 cases only ML experts participated, such as in [LSC*18, NHP*18].

6.6.3. Not Evaluated

As described earlier, approximately half of the papers did not include any type of evaluation. However, we discovered one visualization tool [KTC*19] that was evaluated later in a new publication [KC19].

6.7. Trust Levels

The most novel components of the categorization presented in this section are the different levels of trust we identified in the 200 individual papers. In Section 2, we divided the enhancement of trust in ML models with the help of visualizations into five levels (raw and processed data, learning method, concrete model, and user expectation).

6.7.1. Raw Data (TL1)

Source reliability often comes together with *transparent collection processes* as in AnchorViz [CSV*18], an interactive visualization that facilitates erroneous regions detection through semantic data exploration. By pinning anchors on top of the visualization, users create a topology to lie upon data instances based on their relation to those nearby anchors. Examination of discrepancies between semantically related data points is another functionality of the tool. This data exploration helps to observe source reliability and if any strange effects occurred when the collection process happened. However, as can be seen from the data in Table 3, these two categories are covered rarely by visualization tools.

6.7.2. Processed Data (TL2)

Uncertainty awareness and investigation is an established subject of research in the visualization community with techniques such as the one presented by Berger et al. [BPF11]. The authors developed techniques that guide the user to potentially interesting parameter areas and visualize the intrinsic uncertainty of predictions in 2D scatterplots and parallel coordinates. FairVis [CEH*19] is a recent paper that addresses a new problem which seems to become a trend. *Data bias and equality* is a major issue and should be—as much as possible—removed from our ML models. FairVis facilitates users to review the fairness of ML models in interesting, explored subgroups. iVisClustering [LKC*12] is one of the many visualization tools that allow the *comparison of different structures* (clusters in this case) and *guide/recommend* the users by proposing new clusters based on previous actions. With the help of such visualizations, users can interactively refine clustering results in various ways. Also, iVisClustering can fade away noisy data and re-cluster the data accordingly to demonstrate a meaning representation. Zhao et al. [ZCW*19] developed a tool that enables users to recognize, explain, and choose outliers discovered by various algorithms. Roughly one third of the papers cover outlier detection related topics (68 out of 200), such as [LWT*15, RFFT17, SKB*18].

6.7.3. Learning Method (TL3)

The work of Olah et al. [OSJ*18] tries to *familiarize* users with different DL algorithms. They test robust interfaces that appear when users appropriately combine them and the rich composition of this combinatorial solution space. In our STAR,

interpretability and explainability is separated into four categories: *understanding/explanation*, *debugging/diagnosis*, *refinement/steering*, and *comparison*. These categories may even occur in pairs or triplets, depending on the visualization system and technique. For instance, Liu et al. [LLS*18] support *understanding/explanation* of the reasons behind faulty predictions introduced by adversarial attack examples. The basic concept is to analyze groups of critical neurons and their connections of the adversarial attacks and match them with those of the normal cases. DeepTracker [LCJ*18] facilitates the exploration of the intense dynamics of CNN training processes and helps to identify the unique patterns that are “buried” inside the enormous amount of information in a training log (*debugging process*). Hamid et al. [HDK*19] describe a visual ensemble analysis based on hyper-parameter space and performance visualizations. These visualizations are mutually used with associations’ explorations between topological arrangements and allow the production of enough knowledge in order to support users *steering* the process. Rieck and Leitte [RL15a] suggest a *comparative* analysis of DR methods according to what level of preservation of structural features in the high-dimensional space remains in the 2D embeddings. Local and global structural features are assessed in the original space, and specific DR methods are chosen based on those findings. Manifold [ZWM*19] is conceived of a generic framework that does not rely on the internals of particular ML models and only observes the input and the output. With the comparison of various models and learning methods, it allows users to become *knowledgeable* about their usability. FairSight [AL20] which was discussed earlier along with the What-If Tool [WPB*20] are both two recent examples of how *fairness* is a trending unexplored subject in the community. The What-If Tool (which was not discussed yet) enables domain experts to assess the performance of models in hypothetical scenarios, analyze the significance of several data features, and visualize model functionality across many ML models and batches of input data. It also engages practitioners in grading systems that are able to show multiple ML fairness validation metrics.

6.7.4. Concrete Model (TL4)

Cashman et al. [CPCS20] researched the rapid exploration of model architectures and parameters. To this end, they developed a VA tool that allows a model developer to discover a DL model immediately via exploration as well as rapid deployment and examination of NN architectures. By visually comparing models, beginners might come to similar conclusions (e.g., that early stages of convolutional layers perform well in feature extraction) as ML experts who take advantage of their *experience*. *In situ comparison*, i.e., a comparison of two or more states of the same model, is performed by Gamut [HHC*19], for example. The benefit of Gamut lies in the justification of why and how professional data scientists interpret models and what they look for when comparing their internal components. Our investigation showed that interpretability is not a monolithic concept: data scientists have different reasons to interpret models and tailor explanations for specific audiences, often balancing competing concerns of simplicity and completeness. Moreover, *performance* is one of the most common techniques to choose from (see Table 3) when having different models. LoVis [ZWRH14] allows the user to progressively

construct and validate models that promote local pattern discovery and summarization based on “complementarity”, “diversity”, and “representativity” of models. *What-if hypotheses* are supported by Clustrophile 2 [CD19], which guides users in a clustering-based exploratory analysis. It also adapts incoming user feedback to improve user recommendations, helps the interpretation of clusters, and supports the rationalization of differences between clusterings. Last but not least, papers that deal with issues related to *model bias* and *model variance* usually occur together. Mühlbacher and Piringer [MP13] present a framework for building regression models addressing these limitations. Analyzing prediction bias with model residuals is one of the techniques used to limit the local prediction bias of a model, i.e., avoiding the inclination towards underestimation or overestimation. They also visualize the point-wise variance of the predictions by using a pixel-based view.

6.7.5. Evaluation/User Expectation (TL5)

The *Agreement of colleagues* is related to provenance and the possibility to enable users to collaborate with each other. Wongsuphasawat et al. [WSW*18] present a design study of the TensorFlow Graph Visualizer, which is a module of the shareable TensorFlow platform. This tool improves users’ understanding of complicated ML architectures by visualizing data-flow graphs. These flows can be investigated, and at each point in time, provenance can be considered as a way to return back to a previous situation. *Visualization evaluation*, as mentioned earlier in Sect. 6.6, is activated when the visualization techniques and tools are evaluated or if any type of feedback is provided. Showing *metrics validation/results* is the most common way of enhancing trust until now. Squares [RAL*17] is a performance visualization for multi-class classification problems. Squares supports estimating standard *performance metrics* while demonstrating instance-based distribution information essential for supporting domain experts in prioritizing efforts. Furthermore, Fujiwara et al. [FKM20] implemented a VA method that highlights the crucial dimensions of a cluster in a DR result. To obtain the important dimensions, they introduce an improved method of contrastive PCA. The method utilized is called ccPCA (contrasting clusters in PCA) and can compute each dimension’s relevant contribution to one versus other clusters. An example that implicitly checks *user bias* is the explAIner tool by Spinner et al. [SSSEA20]. explAIner is a VA system based on a framework that connects an iterative explainable ML pipeline with 8 global observing and refinement mechanisms, including “quality monitoring”, “provenance tracking”, or “trust building”. Additionally, Jentner et al. [JSS*18] propose the metaphorical narrative methodology to translate mental models of the involved modeling and domain experts to machine commands and vice versa. The authors provide a human-machine interface and discuss crucial features, characteristics, and pitfalls of their approach. With regard to user bias, the research community has taken “small steps” with only a few papers tackling this issue. However, explicit reports about this challenge are still rare, unfortunately.

6.8. Target Group

In most cases, the visualization tools cover at least the *target group of domain experts/practitioners* [EGG*12, FMH16, FCS*20,

GCRM08, HNH*12, KPN16]. Then, other target groups such as *ML experts* [JC17b, KJR*18, SSK10, WLN*17] and *developers* are in the focus of the authors [KFC16, Mad19, RL15b, YZR*18] (commonly together). *Beginners/novice users* [JSO19, MXQR19, SRM*15, TLRB18] are rarely considered. To give two examples, Bögl et al. [BAF*14] support with TiMoVA-Predict several types of predictions with a holistic VA approach that focuses on domain experts. Providing different prediction capabilities allows for assessing the predictions during the model selection process via an interactive visual environment. Biologists and doctors, for instance, are interested in being able to compare data structures and receive guidance on where to focus on. Ma et al. [MXLM20] employ a multi-faceted visualization schema intended to aid the analysis of ML experts for the domain of adversarial attacks.

7. Survey Data Analysis

In the previous parts of our report, we explained our overall methodology, provided high-level statistical information on the selected papers, and introduced our categorization together with example papers assigned to the individual categories. Now in this section, we discuss lower-level analytical results based on the collected papers and their metadata. In order to detect interesting connections and important emerging topics among the 200 papers, we applied topic modeling to all of them, following the visual text analysis approach by Kucher et al. [KMK18]. While the topic modeling results might be subject to the algorithm and parameter choice concerns to some extent, they provide information complementary to the results of our manual investigations. Thus, the topic analysis contributes both to the validation and the new insights regarding the categorized publications. In addition, we investigate the relations between categories in general (again following the workflow proposed by Kucher et al. [KPK18]) and explore the different data sets used in the individual papers. All those analyses help us to validate and further explore our categorization by creating new insights that can be used as research opportunities for this subject (cf. Section 8).

7.1. Topic Analysis

Methodology. First, we collected the PDF files of the selected papers and converted them to plain text. After that step, we prepared the text corpus by clearing the full texts from the authorship details and acknowledgments. Next, we processed them with the latent Dirichlet allocation (LDA) algorithm [BNJ03, GS04] (a common approach for topic modeling). In order to verify the LDA results—because it might produce diverse results at different executions—we ran the same process several times to get comparable results. The results do not indicate a major deviation from the main topic of each paper previously assigned by the manual categorization process. Finally, our LDA results led to ten topics (limited by us due to the lack of space and our attempt to choose a reasonable number of topics). The top eight terms for each topic are displayed in Table 4 together with the papers belonging to a topic (see S6 for further details). From the terms that occurred, we removed any terms related to the structure of the analyzed texts and not to the actual content, for instance, “figure” and “fig”. Our implementation is based on Python with NLTK [Bir06] for the pre-processing

Table 4: For each of the ten topics, we present the top eight terms that we extracted from the results of the latent Dirichlet allocation (LDA) that has been applied to all papers. The suggested topic titles are shown in italics. Each topic is encoded by one specific color. In each topic, we cite the papers that mostly belong to them.

Topic 1 <i>NNs hidden states and models' parameter spaces visualization (in time series applications)</i>		Topic 2 <i>network visualizations to investigate the behavior of projections and explore classes (for topic analysis)</i>		Topic 3 <i>NNs models' hyper-parameters and reward visualization during training for image applications</i>		Topic 4 <i>vector space, samples, and distributions visualization in projections and DR</i>		Topic 5 <i>models' predictions visualization in clustering (for text applications)</i>	
keywords	papers	keywords	papers	keywords	papers	keywords	papers	keywords	papers
model	[BAF*14, BAF*15, EASKC18, XCH*16, MCZ*17, MBP*11]	topic	[BWZ*20, CAA*19, CWS*17, Glei*13, KKZE20, TLRB18, VKA*18, WSW*18, vv11]	image	[AMJ18, SPBA19, SSS*18, WGZ*19, YCN*15]	projection	[ACD*15, AHH*14, CCZ*16, CD18b, ESS*12, GHRW08, GSS*20, IM*10, JSR*19, KTC*19, LBT*18, LjLH19, SDMT16, SZL*18, WMI*18]	model	[AYMW11, BDSF17, BFFG11, BULBS, CD18a, CD19, DCS09, CWGW19, FMH16, FSJ13, GS14, HNH*12, HVP*19, JCI7b, KEY*18, KKB19, KPB14, KPN16, LKC*12, LRL*18, MP13, MW10, PBK10, RIG19, SBE*18, SKB*18, SML17, SNMM18, SRG*18, ZKM*19, ZWLC19, ZWRH14]
view	13 papers	node	9 papers	model	5 papers	dimension	15 papers	cluster	32 papers
tree		network		training		view		clustering	
state		projection		input		sample		value	
hidden		behavior		network		vector		view	
hidden state		class		hyper		space		prediction	
time		edge		parameter		point		document	
parameter		tree		reward		distribution		variable	
Topic 6 <i>ML models' explanations and visualization systems evaluation</i>		Topic 7 <i>subspaces exploration and distances examination in clustering and DR</i>		Topic 8 <i>models' predictions and design prototyping in ML algorithms</i>		Topic 9 <i>individual points, projection space, and outliers' exploration</i>		Topic 10 <i>NNs neurons' activations visualization during training for image applications</i>	
keywords	papers	keywords	papers	keywords	papers	keywords	papers	keywords	papers
model	[ASW*19, BAL*15, BEF17, CH*19, CHH*19, CS14, DCC*19, GHC*19, GSC16, GSK*20, HHC*19, HSD19, KPB16, JKM12, JSO19, JSS*18, JZF*09, KBWS15, KDS*17, KFC16, KJR*18, KLTH10, KPB18, KSH19, MMD*19, MSM*17, MSW10, PSF17, RAL*17, SKKC19, SLT17, SSK10, SSS*20, TKOB17, TWF*12, TSL*16, WML*19, WPB*20, ZYB*16]	cluster	[AEM*11, BAPB*16, SLG*12, OGB*19, CHAS18, CLKP10, CSG*18, FCS*20, FKM20, JHB*17, JJ09, KDFB16, LMZ*14, LWBP14, LWT*15, ML14, MWM11, PSM12, RL14, RL15a, RL15b, RL16, TPRH11, VLO9, WLN*17, WLS19, XVC*18, ZYR*16, ZLH*16, dBD*12]	model	[CPCS20, CSV*18, GDM*19, KAKC18, KCK*19, LGS*18, LLL*19, MGB19, MXC*20, MXLM20, MXQR19, SGB*19, SJS*17, SJS*18, SPG14, WGSY19, XXM*19, ZWM*19]	point	[AL20, AuP*07, BHR*19, S2L*18, CMN*16, RSF*15, JPN15, KKW*17, KS12, LA11, MCM14, PLHL19, RFT18, SNLH09, SRM*15, SMLB10, YMT17, ZCW*19, ZWC*18]	layer	[AJY*18, CGR*17, GFM*18, HKB*19, HLW*20, HPRC20, LCG*18, LLS*18, LSC*18, LSL*17, LXL*18, NHP*18, OS*18, PHV*18, PSMD14, RFFT17, SW17, WGY18, ZF14, ZHP*17]
learning		subspace		prediction		class		network	
system		dimension		instance	18 papers	projection		neuron	
machine		point		view		space		model	
participant		view		prototype		view		training	
explanation		distance		algorithm		instance		activation	
instance		dimensional	30 papers	learning		value		class	
machine learning	39 papers	clustering		design		outlier		image	

of stop words and Gensim [RS10] for the topic modeling part. The names of the topics were manually assigned by us after several discussion cycles considering both the top terms and the contents of the papers in each topic. The results are then visualized with the assistance from the interactive visualization tool described by Kucher et al. [KMK18], see Figure 5. This visualization is based on a DR projection which may not be the most reliable approach. However, the ground truth labels taken from the LDA results match in almost all the cases with the clusters formed by the embedding.

Topics. In the following description list, we shortly summarize the ten topics (see Table 4) we identified:

■ **Topic 1 – hidden states & parameter spaces.** According to our analysis, the common factor between the majority of the 13 papers in this topic class is their focus on time series data [BAF*14, BAF*15] and RNNs [SBP19] (in Strobel et al. [SGPR18]: long short-term memory networks). It seems that the exploration of the hidden states of such networks preserves lost information that could enhance trust [MCZ*17] with appropriate expert intervention. Another subtopic here is the ML models' parameter spaces exploration [TCE*19], which enables users to find the best parameters based on a series of optimizations for certain goals. In this context, Mühlbacher et al. [MLMP18] present an approach that visualizes the effects of these parameters. As stated by these previous works, support for the visual parameter search is still an open research challenge.

■ **Topic 2 – investigation of the behavior.** This topic class contains 2 out of 9 papers on topic analysis applications [CAA*19, KKZE20]. A common theme here is network visualization used for explaining Bayesian networks [CWS*17, VKA*18] and decision trees [vv11]. Other subtopics (which lead to research opportunities) are the exploration of behavior with regard to the decomposition of projections, showing the internal parts of ML models (and how classes are formed inside them), and the role of the user; they are all covered by our categorization presented in Section 6.

■ **Topic 3 – hyper-parameters & reward.** All five papers in this class (except one [SPBA19] related to reinforcement learning) make use of image data. They form a tight, green cluster in Figure 5 and focus mainly on DL techniques [YCN*15]. It is an open challenge to decide which hyper-parameter value [SSS*18, WGZ*19] is better for a particular NN and specific problem, thus implying novel visualization techniques. For reinforcement learning, instead, more research is needed for studying what behaviors are associated with the two types of reward (high and low) and monitoring how they develop during training.

■ **Topic 4 – vector space, samples, & distributions.** In this topic class, three papers [LBT*18, LjLH19, SZL*18] work with vector space embeddings that illustrate similarities of the data. Most of the 15 papers highlight the importance of visualizing instances and samples [KTC*19, SDMT16] that form clusters in projections and explore the distributions [AHH*14, CCZ*16, CD18b] of points in DR techniques.

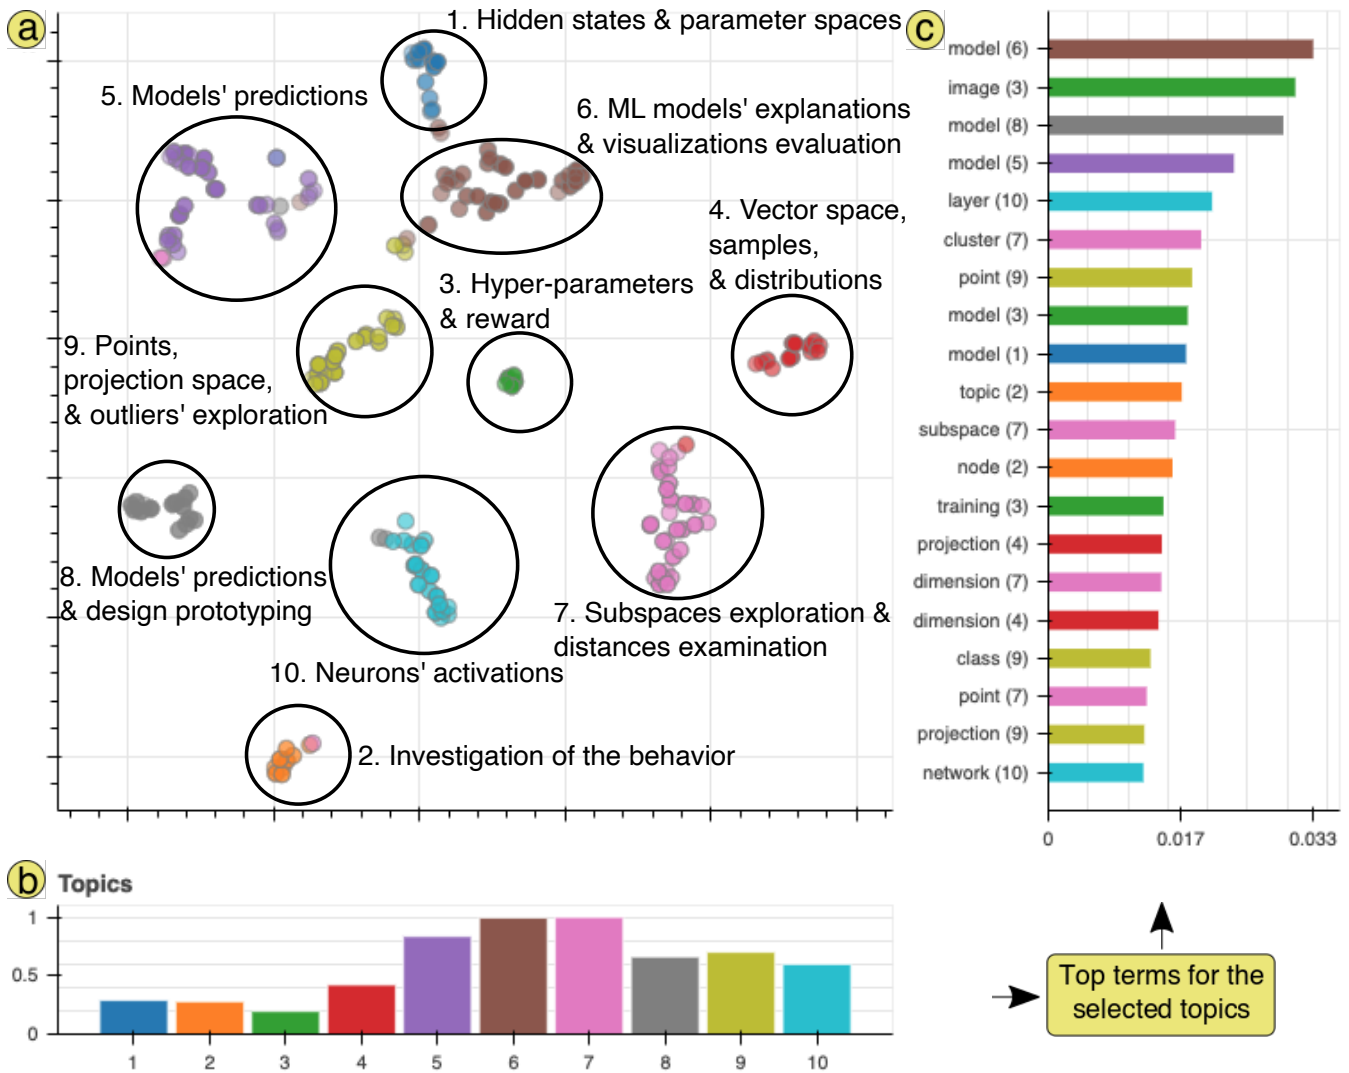


Figure 5: Visual exploration of new interesting topics derived from the 200 papers. (a) Papers' embedding generated with the t-SNE algorithm and based on the corresponding topics. The black outlines were manually drawn on top, and the tags act as short versions of the 10 full topic titles. (b) Bar chart of topics with each topic's significance (scaled from 0 to 1). (c) Horizontal bar chart of top terms with the highest relevance for all the topics. Here, topics are encoded with color and number (in parentheses); a single term can be found in several topics.

Finding ways to improve these visualizations is still an open challenge in the InfoVis community.

Topic 5 – models' predictions. Models' predictions and results visualization with the use of *quality and validation metrics* [FSJ13,GS14] (depending on the ML type) composes a big, more general topic class with 32 surveyed papers. A subgroup in this class especially refers to clustering challenges [BDSF17, KEV*18] and open research questions such as: “do we have the best clustering that could be achieved and if not, how can we improve it?” (usually related to text applications) [CDS09].

Topic 6 – models' explanations & visualization evaluation. This topic is rather generic (with 39 papers allocated) as it addresses the *understanding/explanation of ML models* [GHG*19,TKDB17]. For many visualization tools belonging

to this topic class, we can observe that *user studies* (i.e., evaluations) [BAL*15, BEF17, GSC16, KLTH10, SLT17, ZYB*16] have been performed with participants from different educational levels (novices, practitioners, ML experts, and so on). Following the overall theme of this STAR, a straightforward unsolved problem in this area is to find answers to how exactly we shall progress with the development of visualization tools for boosting the trust in ML models and their results.

Topic 7 – subspaces exploration & distances examination. Clustering and DR are both covered together when exploring subspaces [BAPB*16, KDFB16, LMZ*14]. Finding the correct distance function, checking if these distances are preserved after the projection from the high-dimensional space into the 2D space, and matching the *users' cognitive expectations* is clearly

not a trivial task. As a result, many papers are published in this area [AEM11, BLBC12, JHB*17, WLS19], making this topic class with 30 papers one of the most prominent in our analysis.

■ Topic 8 – models’ predictions & design prototyping.

Another generic topic class with 18 related papers contains, among others, the subject of ML models’ predictions [SJS*17, XXM*19] that has already been seen in Topic 5. The difference between this class and Topic 5 is the focus of its related papers, which is on the instantiation of visualization prototypes with different design choices that should be carefully considered based on previous InfoVis research. As such, updating the current methods with improved versions can lead to enhanced trust of visualizations and *reduce biases* [GDM*19, LGG*18, SGB*19].

■ Topic 9 – points, projection space, & outliers’ exploration.

With 68 papers in total, the area of *outlier detection* is prominent in our categorization, see Table 3. This category is even confirmed through the topic analysis (with 19 papers assigned) as many techniques work with *outlier detection* [BHR*19, JPN15, RSF*15]. Another hot topic in the visualization community is the visual analysis of relations between points and dimensions within the various projection spaces [Aup07, CMN*16].

■ **Topic 10 – neurons’ activations.** When visualizing DL techniques, the existing research has tried to address the activation of neurons in NN and their visual representation [HDK*19, HLW*20, HPRC20]. Different visualization techniques (e.g., 2D saliency/activation maps) have been used to visualize the activations of such neurons in various DL models, especially for image applications [AJY*18]. This topic class consists of 20 papers about visualizing the internal operations of NNs during the training phase. A possible research question in this context is: “what else can be visualized (for instance, gradients [CPM*18]) that gives meaning to humans about the learning process of a NN?”

Topic embedding. The ten-dimensional data space of the topics over all 200 papers has been reduced to two dimensions by using t-SNE [vdMH08], i.e., two papers are positioned close to each other if their topic relationships are alike, see Figure 5(a). The scales in the depicted bar charts are from 0 to 1, with 1 being the highest relevancy value of a topic in Figure 5(b) and of a term in Figure 5(c). The black outlines in the 2D embedding (see Figure 5(a)) were appended manually.

As it can be derived from Figure 5(b), Topics 6 & 7 are the most prominent ones, followed by Topics 5 & 9, 8 & 10, and the others. In more detail, *ML models’ explanation and visualization systems evaluation* (Topic 6) and *subspaces exploration and distances examination in clustering and DR* (Topic 7) are two discussed topics that cover approximately 35% of all papers. With regard to Figure 5(c), some interesting top terms are—as expected—“models” (ML), “image data” (computer vision, see even Table 5), “layers” (DL), “clusters”, “topic” (analysis), “subspace”, “projections”, and “dimensions” (DR). By observing the t-SNE projection in Figure 5(a), we can find more interesting insights. For instance, the tightest cluster is color-encoded in green and related to *NNs models’ hyper-parameters and reward visualization during training for image applications* (Topic 3). Another interesting result is that the misclassification of orange (Topic 2) & pink (Topic 7)

points as well as of pink (Topic 7) & red (Topic 4) points in the embedding happens due to three concept terms that are spread in all three topics, namely, the terms “clustering”, “dimension”, and “projections”. Furthermore, as Topic 6 is rather generic (*ML models’ explanations*), there are some points laid out in-between (i.e., mixed points) with Topics 1 (DL) & 9 (*projections*). Lastly, Topic 8 (*models’ predictions & design prototyping*) is also rather general, because the points in the projection are spread through two other topics (5 & 10); and this is probably because NNs are a subclass of ML models and Topic 5 (*models’ predictions*) is very similar to Topic 8.

Overall, we notice that the automatically generated topics introduce new subcategories (and ideas) that have been discussed in parallel to our categorization and—in consequence—supported even more the categorization of the papers described in the previous section. For instance, Topics 1 and 10 represent VA tools focusing on the visualization of the NNs hidden states and neurons’ activations respectively to facilitate the *understanding/explanation* of them. In addition, Topic 1 covers examples for the *comparison* of models based on the visualization of their parameter spaces. Similarly, Topic 2 is related to the *in situ comparison* of concrete models to investigate different behaviors of the ML models. Topic 3, instead, focuses more on *diagnosing/debugging* the training process for reinforcement learning, and Topic 4 reflects the *comparison of data structures* with the use of projections and DR. The remaining topics are explicitly connected to our TL categorization within the corresponding topic list items above. We believe that this mixture of coarse-grained manual categorization with a fine-grained automatic refinement may help guiding potential readers to more insights and analyze the surveyed papers even further.

7.2. Correlation and Summarization of Categories

Correlation between categories. We have conducted a correlation analysis for the categories used in our collected survey data set. Individual visualization papers were treated as observations, and categories (cf. Table 3 and S5) were treated as dimensions/variables. Linear correlation analysis was then used to measure the association between pairs of categories. The resulting matrix in Figure 6 contains Pearson’s r coefficient values and reveals specific patterns and intriguing cases of positive (green) and negative (red) correlation between categories. Since the interpretation of the coefficient values seems to differ in the literature [Coh88, Eva96, Tay90], we focus on values of correlations that appear interesting to us despite a potentially strong or weak correlation level. Due to the extensive size of the correlation matrix, we include only a thumbnail of it and refer the reader to S7 for more detail. In Figure 6, we present some strong, medium, and weak correlation cases that caught our attention.

The strongest case of *negative* correlation in our data set is the *not evaluated category vs. user expectation for evaluation* (cf. 6.6. and 6.7.5.), which clearly highlights the need for further evaluation of visualization tools and techniques. Further interesting cases mainly include competing categories from the same group. For example, *model-agnostic* techniques contradict *model-specific* techniques, because they consider different visualization granularities for a given ML model. 2D and 3D oppose each other as typically

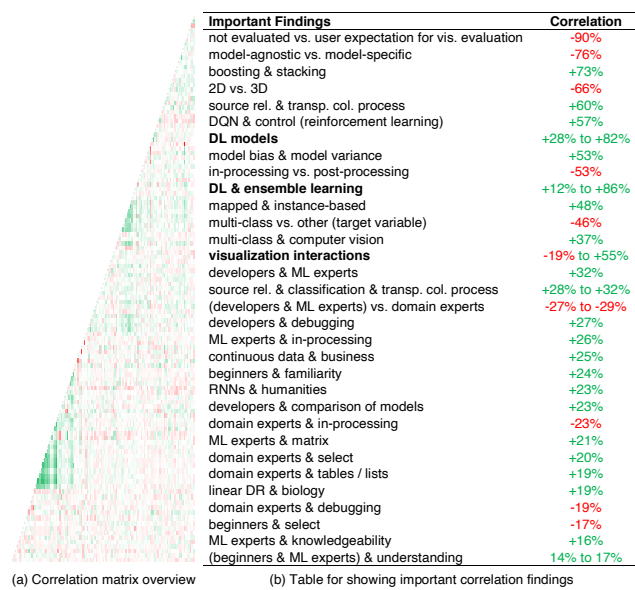


Figure 6: The matrix of correlation values for categories (see Table 3) calculated from our survey data set. In subfigure (a), we show the correlations for the categorization in a thumbnail version of the correlation matrix. The entire correlation matrix for further exploration can be found in S7. (b) To help the reader to see potentially interesting findings, we list the cases in decreasing order of their correlation strengths (independent of negative or positive correlation). In both subfigures, we use red color for negative correlation and green for positive. Cases highlighted in bold refer to groups of categories in general (e.g., visualization interactions).

only one of them exists in a visualization approach. Moreover, techniques that focus on data exploration, explanation, and manipulation related to the *in-processing* phases of an ML pipeline are very different compared to systems that monitor the results in the *post-processing* phase of an ML model. The strong negative correlation between *multi-class* and *other* target variables might point to an effect that comes from our own categorization procedure: when papers could not be mapped to a concrete target variable (*multi-class*, for instance), then the *other* category has been assigned, e.g., to show the irrelevance of the target variable for a visualization technique. The category *domain experts* is negatively correlated to managing models during the *in-processing* ML phase, which makes sense as they do often not know much about how models work. Similarly, *developers* and *ML experts* together are weakly but negatively correlated with *domain experts* confirming the previous acquisition. Other insights are that *beginners* are not usually using *selection* as interaction technique and *domain experts* do not work with *diagnosing/debugging* ML models as they do not have the experience and/or knowledge following the previous inference.

Cases with *positive* correlation start in Figure 6(b) with *stacking* which is highly correlated with *boosting* ensemble learning as the former sometimes includes the latter technique. All *DL techniques* among each other have on average a medium positive correlation, which shows that they have much more in common com-

pared to other ML methods, for example, DR. The same is true for the group of *visualization interactions* to a slightly less extent. When *source reliability* is taken into account and researched by scientists, then the *transparent collection process* is usually examined together. *Deep Q-networks (DQNs)* are positively correlated to and seem to be normally used together with *reinforcement learning* methods, particularly with the subtype *control*. Furthermore, when *model bias* challenges are addressed by visualization, then *model variance* is another category that is addressed simultaneously. *Ensemble learning* along with *DL* are positively correlated, as the former includes the latter in many cases (a fact already mentioned before). *Mapped instances* lead to *instance-based* visualizations, in general. With regard to domains, *continuous* target variables and *business* are positively correlated (potentially due to trend predictions); as well as *computer vision* and *multi-class* data in an even stronger fashion. The latter correlation is also supported by our data set analysis (see Table 5). Finally, *visualization interactions* are positively correlated among each other, with the exception of *verbalization* which is negatively correlated with the remaining categories in this group. This possibly means that *verbalization* is not frequently used by the VA system developers. For more details, we refer the reader to Figure 6(b) and supplementary material S7.

Popular approaches. The statistics in Table 3 support our expectations of the most common aspects of existing visualization techniques for enhancing the trustworthiness of ML models. For our first aspect (6.1. Data), *computer vision*, *humanities*, *health*, and *biology* seem the most prominent *domains* in the surveyed papers. *Multi-class classification* is the most common target variable in our discussed techniques. Furthermore (6.2. ML), *linear* and then *non-linear DR* techniques are commonly used, followed by *bagging (ensemble learning)* and *CNNs* from the *DL* class. The vast majority of the papers address *supervised learning* and specifically *classification* problems, and in second position *DR* and *clustering* which belong to *unsupervised learning*. (6.3. and 6.4.) *Post-processing* and *model-agnostic* visualization techniques cover around 75% of all papers. (6.5.) With regard to *visual aspects* and *granularity*, almost all techniques used have at least a component which is *computed* and not *mapped/derived* from the data directly; and *aggregated information* is slightly more common than *instance-based/individual* exploration of instances.

The absolute majority of the visualizations rely only on *2D representations*, and *color* is the visual channel most commonly used for encoding information in the corresponding visualization systems, tools, and techniques. The rather large number of techniques using *opacity* to hide points/instances and *size/area* to encode data attributes can be explained by the extensive usage of *scatterplots*. Other popular visualizations are *bar charts*, *custom glyphs* and *specialized icons*, *histograms*, and finally, *line charts*. More traditional visual representations, such as *tables*, *lists*, and *matrices*, are working in pairs with *instance-based* exploration techniques, which are far less complicated than the previously mentioned visualizations. On the interaction side, *selection*, *exploration*, and *abstraction/elaboration* are the three most prominent categories found in many papers, followed by other interaction techniques, such as *connecting* all the different views, *filtering out* or *search-*

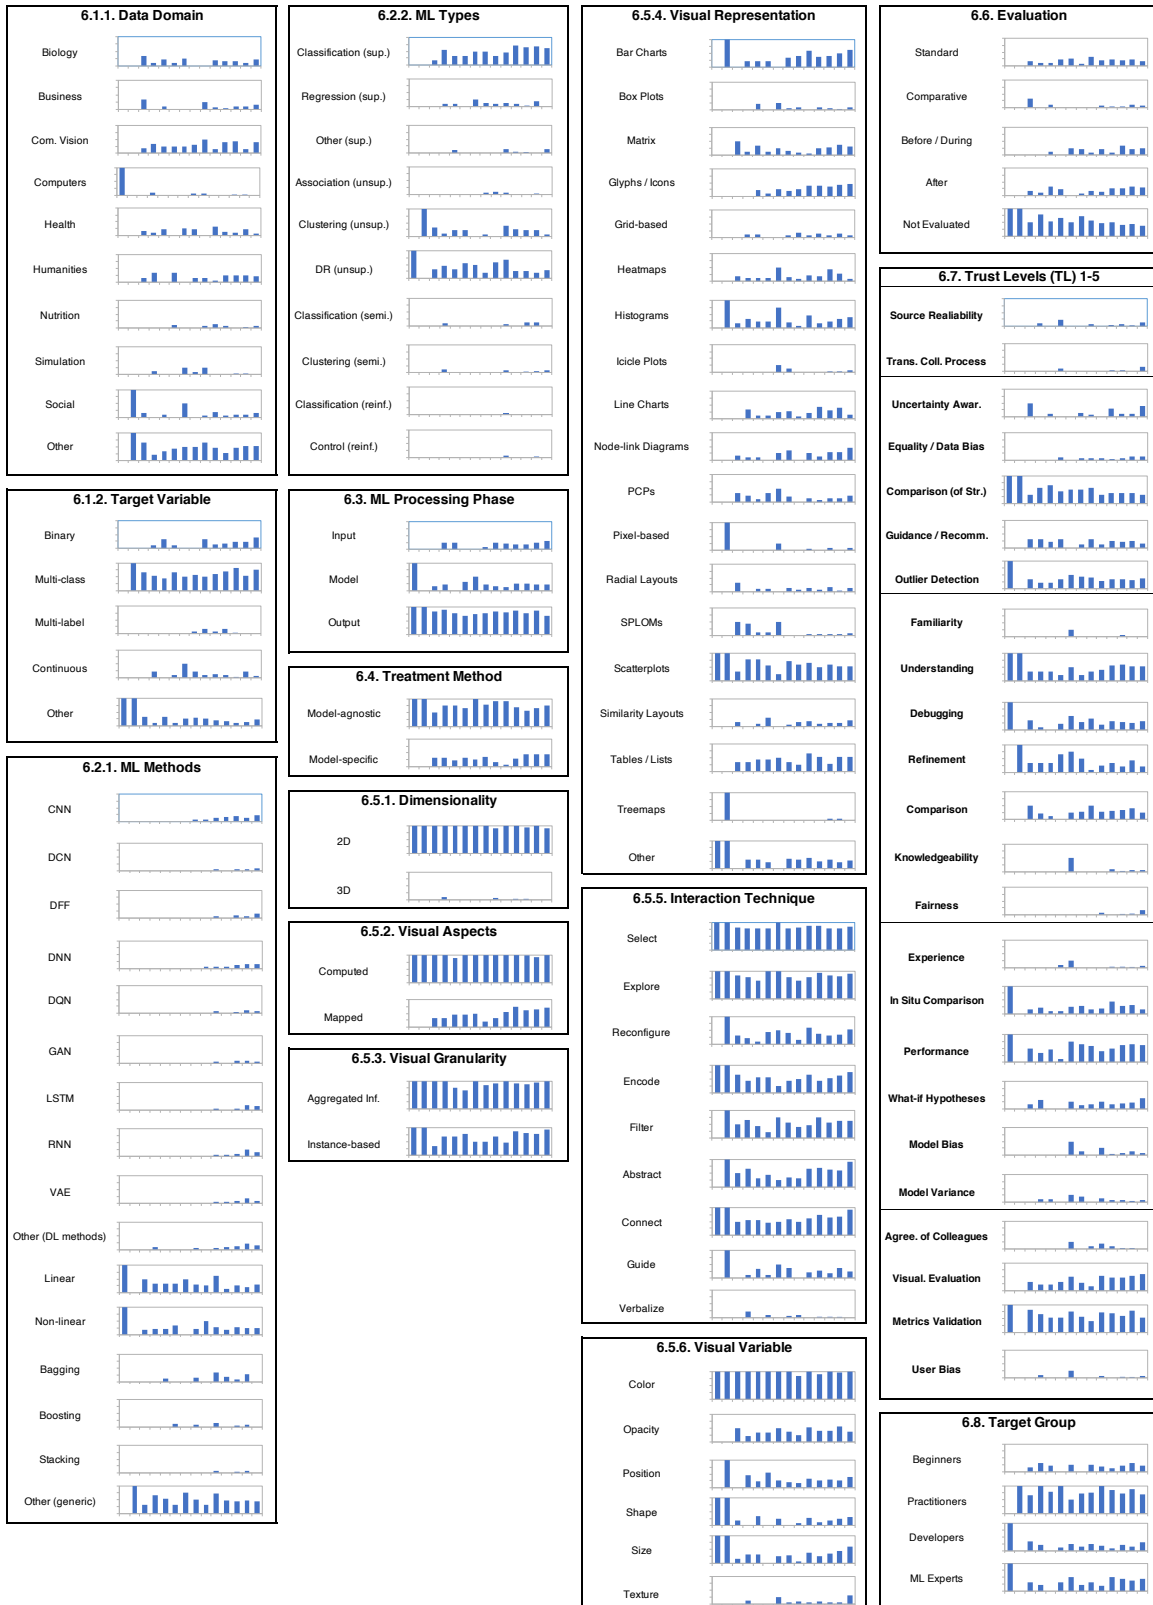


Figure 7: Sparkline representation of relative category popularity over time during 2007–2020. Each bar represents the support for a specific category (cf. Table 3) in relation to the total count of techniques in the same year in our data set.

ing for specific instances, and *encoding*. (6.6.) Around half of the visualization techniques that we analyzed have *not been evaluated*.

The *trust levels* (6.7.) show that more works tackle *source reliability* problems rather than the *transparent collection process* challenge (as seen in TL1). For the second level (TL2), researchers focus on the *comparison of structures* and *outlier detection*. In the third level (TL3), *understanding*, *steering*, *comparing*, and *debugging* ML methods are quite popular. These aforementioned categories can be considered under the umbrella of interpretable/explainable ML methods. For TL4, *performance*, *in situ comparison*, and *what-if hypotheses* are other very often occurring categories connected with the selection process of an individual ML model. Ultimately in TL5, *metrics validation and results observation* at the final stage of the processing phase is the most frequent category with 130 papers. Last but not least for 6.8., the visualization systems and techniques have as a main *target group*, usually *practitioners/domain experts*, followed by *ML experts* with large distance. The analysis above sheds light into the reasons why a few approaches seem to be more popular than others. The ML side uses mostly performance, metrics validation, and results to monitor and boost trust in the ML models. In contrast, the visualization side focuses more on traditional visual representations and/or multivariate, scalable visualizations that the experts are more willing to use.

Temporal trends. While the analyses presented above focus on the overall statistics, we have also analyzed the temporal trends for individual categories based on the collected data. Figure 7 provides a sparkline-style representation of the information about each category's support (i.e., the count of corresponding techniques) over time. The values are normalized by the total count of techniques for each respective year between 2007 and 2020 (for example, 3 out of 9 papers from 2010 used *computer vision data* to demonstrate the usability of their tools). The resulting representation in Figure 7 allows us to confirm, for instance, that the *ML processing phase* visualized consistently most often is *post-processing* rather than *in-processing* or *pre-processing*. Combination of such temporal trends with the overall statistics also allows us to identify and further discuss the usage of currently underrepresented categories.

Underrepresented categories. *Multi-label data* and *computer-related data* (from software or hardware) are two underrepresented categories that show no trend for a potential increase according to Figure 7. For *ML methods*, approaches such as *stacking ensemble learning*, *deep convolutional networks (DCNs)*, and *DQNs* are also not covered in detail. Nevertheless, there is a very small increasing trend for them observable in Figure 7. Explicit techniques addressing problems that come with *stacking ensemble learning* were not found in any paper, thus indicating a new research opportunity. For *ML types*, the subcategory of solving *classification* problems while using *reinforcement learning* is almost never visualized and actually never addressed explicitly by the visualization community. Other underrepresented categories here are *reinforcement learning* and *control*, and *association for unsupervised learning*.

For the *visual representation*, the *treemaps* and *icicle plots* categories are virtually not supported by the data. Further techniques that belong to the last category within *visual representation* ("other") and are fairly underrepresented are *waterfall charts*,

bipartite visualizations, and lastly *area charts* (as also mentioned in Section 6). For the *interaction techniques*, the category of *verbalization* emerged in 2010 and has not attracted much support in the publications; even though recently in 2018, Sevastjanova et al. [SBE*18] argued about the importance of its existence. Moreover, *texture* is the least usual way to represent the data visually in comparison to the others. *Comparative evaluations* are the rarest way of *evaluating visualizations*, which is rather logical because not every technique has an obvious opposing one.

The real challenges start when we check the *trust levels* aspect because many techniques are underrepresented, which means there are several research opportunities in the area. *Transparent collection processes*, *source reliability*, and *equality/data bias* are usually not covered by papers. Other problems, such as how visualization can assist with the *familiarity* a user has for a *learning method*, should also be in the research agenda of our community. *Fairness* (and previously mentioned *equality* for the data) of the *learning methods* seem to be in the spotlight according to the temporal statistics (see Figure 7). Finally, *developers* (i.e., model builders) and *beginners* are the two most underrepresented target groups in the papers we analyzed. *Knowledgeability about learning methods* and details available to different types of users is not well supported. As a result, customization and reconfiguration of visualizations that take into account the *experience* of users in order to choose a specific ML model are not researched to the required extent. Furthermore, a few techniques enable *agreement of colleagues* and study about the consequences of using provenance in visualization tools in order to cover our discussed subject. *User bias* is ignored in almost all of the visual systems.

All of the underrepresented categories discussed above might be candidates for open challenges, as can be seen in Section 8.2. From an ML perspective, the most real-world challenges are about either classification or regression problems. Consequently, other ML types are not researched to the same level. From the visualization perspective, a large amount of time and effort is necessary to design and perform a "proper" visualization evaluation [LTBS*18]. Moreover, as long as the visualization tools do not focus on beginners, familiarity with and knowledgeability of the algorithms are left aside by visualization researchers.

7.3. Data Set Analysis

Methodology. For the data set analysis, we consider only non-synthetic (i.e., not artificial) data sets which can be accessed online. We also include data sets that can be requested from the paper authors. For the individual data features, we take further into account the labels (i.e., classes), if they are existent. Overall, details about a data set were collected relying on the description provided by the authors of a paper, for example, how they collected and stored the data. In any other case, we omitted the data sets. All data sets are sorted first according to the number of occurrences in the 200 papers, and then by year to show the most recent first.

Results. The result of this process can be observed in Table 5. In the listed 38 cases, the data sets are used in at least two papers, and the remaining 106 entries are used once only (cf. S8). In total, we managed to identify 144 non-synthetic data sets in our

Table 5: Overview of data sets ordered according to their usage. A “#” in the columns indicates the number of papers in this survey using a particular data set; the data sets are then further grouped based on this number.

Data Set	Year	Domain	# Instances	# Features	Target Variable	#	References
MNIST [LBBH98]	1998	Computer Vision	70,000	784 (28x28)	Classification (10 classes)	17	[AMJ18, BZL*18, CHAS18, FKM20, HDK*19, JC17b, LSL*17, MXLM20, PLHL19, PHV*18, RAL*17, SLT17, SSSEA20, WGY518, WGZ*19, XYC*18, CD18b]
Iris [And36]	1936	Biology	150	5	Classification (3 classes)	12	[BZL*18, GSS*20, CHAS18, JHB*17, JZF*09, RL16, SDMT16, SMSL17, SNLH09, WLN*17, WM18, FCS*20]
Wine Quality (red & white) [CCA*09]	2009	Business & Physicochemical	6,497	12	Classification & Regression	10	[BLBC12, CMN*16, FKM20, HHC*19, JZF*09, SMSL17, SNLH09, TLRB18, ZCW*19, SRM*15]
ImageNet [DDS*09]	2009	Computer Vision	Depends (e.g., 1,350,000)	Depends	Classification (1,000 classes)	8	[AJY*18, AMJ18, ASW*19, HPRC20, JC17b, YCN*15, ZF14, LCJ*18]
Food and Nutrition [USD19]	2019	Nutrition	Depends (e.g., 7,637)	Depends (e.g., 14)	Clustering	7	[FKM20, KDFB16, KKW*17, TMF*12, WLN*17, WLS19, ZLH*16]
CIFAR-10 [Kri09]	2009	Computer Vision	60,000	1,024 (32x32)	Classification (10 classes)	6	[CPCS20, HDK*19, JC17b, LSC*18, LSL*17, WSW*18]
20 Newsgroups [Lan95]	1995	News & Humanities	18,828	2	Classification (20 classes)	6	[ASW*19, KBWS15, KLTH10, LKC*12, SLT17, SNMM18]
Yelp Open Data Sets [Yel19]	2019	Reviews & Other	Depends (e.g., 6,685,900)	Depends (e.g., 9)	Clustering, Classification	4	[MCZ*17, MXC*20, MXQR19, TKDB17]
Google's Quick Draw [JRK*16]	2016	Computer Vision	50,000,000	784 (28x28)	Classification (345 classes)	4	[CD18b, CPCS20, WGY518, WGZ*19]
Reuters-21578 [RSW02]	2002	News & Humanities	21,578	5	Classification	4	[LGG*18, KS12, SNMM18, SSK10]
Optical Recognition of Hand Digits [AK98]	1998	Computer Vision	5,620	64	Classification (10 classes)	4	[BAPB*16, BDSF17, KEV*18, LA11]
Census (Income) [DG17, YRW07]	1994	Social	48,842	14	Classification	4	[MBD*11, GNRM08, WPB*20, MLMP18]
Boston Housing Prices [HR78]	1978	Business	506	13	Regression	4	[AMJ18, ASW*19, HHC*19, SNLH09]
Numbeo Quality of Life Index [NUM15]	2015	Socioeconomic	86	7	Clustering	3	[KKW*17, ZWC*18, ZCW*19]
OECD Better Life Index [Lin14]	2014	Socioeconomic	34	24 (11 main)	Clustering	3	[CD18b, CD19, SDMT16]
House Prices [Coc11]	2011	Business	2,919	81	Regression	3	[DCC19, KP18, SRG*18]
Caltech-101 [FFFP04]	2004	Computer Vision	3,030	Depends (e.g., 60,000)	Classification (101 classes)	3	[KLTH10, TLKT09, ZF14]
Communities and Crime [RB02]	2002	Social	1994	128	Regression	3	[FKM20, FSJ13, KDFB16]
Corel Image Features [DG17]	1999	Computer Vision	68,040	89	Clustering	3	[MCMT14, PSPM12, SDMT16]
Image Segmentation [DG17]	1990	Computer Vision	2,310	20	Classification (7 classes)	3	[CMN*16, JPN15, vv11]
COMPAS Recidivism Risk Score [COM19]	2019	Justice	Depends	Depends	Classification (2 classes), Clustering	2	[AMJ18, CEH*19]
USPS Hand. [USP17]	2017	Computer Vision	9,298	256 (16x16)	Classification	2	[LWT*15, ZTR16]
Titanic [Tit15]	2015	Life	712	10	Classification (2 classes)	2	[HHC*19, ZWLC19]
Parkinson's Disease [GTS*08, PPM14]	2014	Health	Depends (e.g., 8,652)	Depends (e.g., 37)	Clustering, Classification, Regression, Clustering	2	[CD18a, CD19]
Tennis Major Tournament [DG17]	2014	Sports	127	42	Regression, Clustering	2	[FKM20, LRL*18]
ALL [PSPM12]	2012	Publications	2,814	1	Clustering	2	[CMN*16, PSPM12]
Software Projects Quality [MSM*10]	2010	Computers	6,773	12	Clustering	2	[CMN*16, SRM*15]
Combustion Simulation [HSPC06]	2006	Simulation	2,800	10	Clustering	2	[LWBP14, WLN*15]
University Ranking [DG17]	1988	Education & Status	285	18	Classification	2	[JHB*17, KEV*18]
Synthetic Faces [TSL00]	2000	Artificial Intelligence	698	4,096 (64x64)	Clustering	2	[RL15a, RL15b]
Concrete Compressive Strength [Yeh98]	1998	Construction	1,030	9	Regression	2	[LMZ*14, RL15b]
Adult [DG17]	1996	Social	48,842	14	Classification	2	[CEH*19, SJS*18]
German Credit [DG17]	1994	Financial	1,000	20	Classification	2	[AL20, ZWLC19]
ISOLET [DG17]	1994	Computers	7,797	617	Classification	2	[Aup07, SvLB10]
Waveform Database Generator [BFOS84]	1984	Telecommunications	5,000	22	Classification (3 classes)	2	[BAPB*16, LMZ*14]
Auto MPG [DG17]	1993	Business	398	9	Regression	2	[XYC*18, WLS19]
Penn Treebank [MMS93]	1993	Humanities	60,774	36	Classification	2	[SGPR18, SBP19]
Breast Cancer Winconsin [BM92]	1992	Health	699	10	Classification (2 classes)	2	[MQB19, XXM*19]
Others (i.e., data sets that are used once)			106 additional unique entries, see supplementary material S8				

200 surveyed papers. The most frequently occurred data sets are MNIST [LBBH98], Iris [And36], Wine Quality [CCA*09], ImageNet [DDS*09], Food and Nutrition [USD19], CIFAR-10 [Kri09], and 20 Newsgroups [Lan95]. 3 out of these 7 data sets are about computer vision and are usually used in papers that work with DL and NNs. Validating our previous categorization, *classification* and then *clustering* problems are the more occurring *target variables*,

and finally *regression*. The number of instances and features can be found in our table along with the number of classes for some cases (if available). The individual papers that used the data sets are listed in the rightmost column of Table 5; and references to the data set providers are given together with the name of the data sets in the first column.

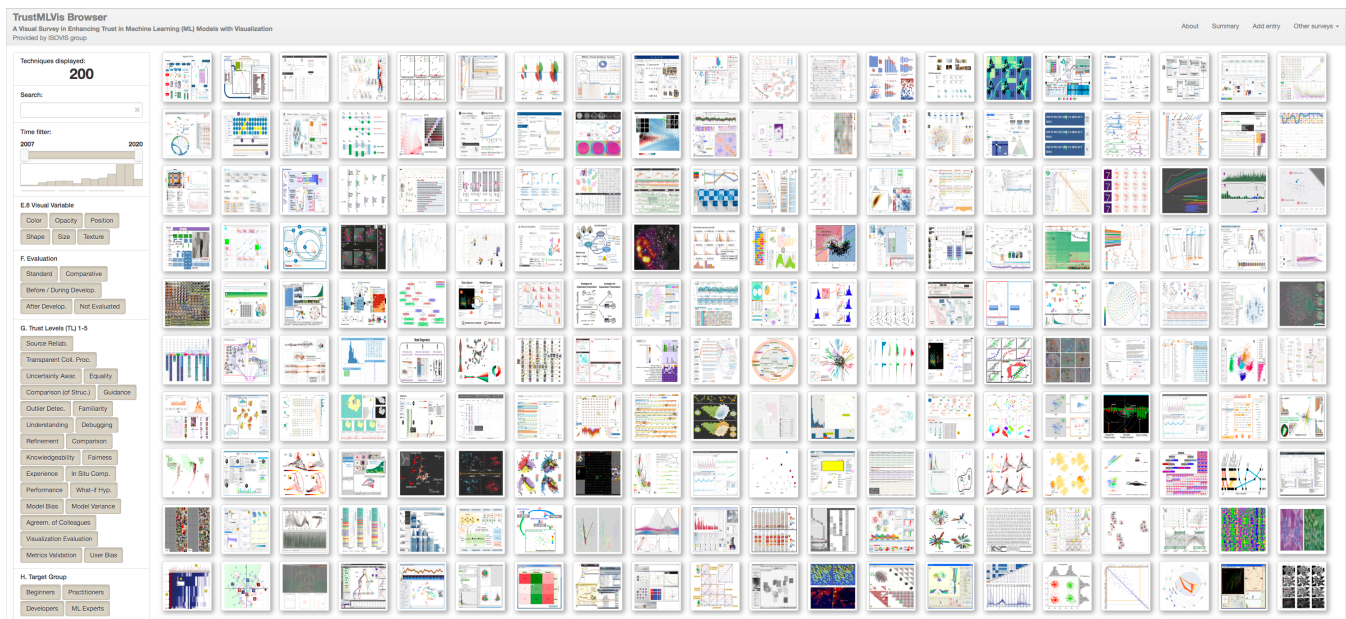


Figure 8: The user interface of TrustMLVis Browser (available at trustmlvis.lnu.se), an interactive survey browser accompanying this article.

8. Discussion and Research Opportunities

In this section, we discuss our online survey browser. Afterwards, we move on to research opportunities based on the data-driven analyses presented in Section 7.

8.1. Interactive exploration with a survey browser

Our work on this survey has been complemented by the development of an interactive survey browser [BKW16, KK14, KK15, KPK18, Sch11, TA13] similar to our group's previous contributions on text and sentiment visualization. TrustMLVis Browser is available as a web application, and its user interface (see Figure 8) comprises (1) a grid of thumbnails representing visualization techniques and (2) an interaction panel supporting category-, time-, and text-based filtering. The user can access the details and bibliographic information about a specific technique by clicking on the corresponding thumbnail. Several dialogs with the overall statistics for the complete data set (cf. Table 3) and the supplementary materials are available via the links at the top of the web page. We encourage the readers of this article to explore the data with the survey browser and to suggest further candidate entries by using the corresponding “Add entry” dialog.

8.2. Research Opportunities

The impact of bias. By looking at our categorization, we can infer that some level of *bias* might be represented in all our defined trust levels in different forms: (a) *data bias (equality)*, (b) *previous familiarity with algorithms*, (c) *model bias*, and (d) *user bias*. Also, it is known that visualization techniques ordinarily do not scale very well when analyzing massive volumes of data. Furthermore, some of the ML approaches have inherent challenges to face, for example, the curse of dimensionality [Bel03] in case of

DR. Thus, considerable levels of selection bias might be unintentionally ignored by the user, for instance, when users have to choose from a selection while not seeing the entire picture and/or the alternatives [GSC16, LA11]. Hence, the research question here is: “what novel solutions can help users to minimize the impact of bias with regard to the data?” A potential answer would be to consider various interaction logs with the VA system. Data generated as part of the analysis process could be considered as well. This data together with the logs could be processed automatically with additional independent ML models and potentially guide users to improvements of the underlying ML models used in the data analysis process. Hence, the ways of combining automatic methods with smart visualizations [Shn20] are still not revealed and should be further evaluated with empirical studies as well as quantitative and qualitative experiments.

Alternatives and combination. Visualization is often used as the medium enabling human-computer interactions (HCI). It usually encourages the development and application of multidisciplinary methods originating from different areas of research. To find an equilibrium state between human and computer controlling the ML process is not a trivial task [Shn20]. Researchers that are intimate with ML models and visualizations are capable of appropriately promoting the joint development of visual explanations for ML models. Furthermore, there is a possibility to employ verbalization (as discussed before) as a complementary tool alongside visualization for explaining ML models. The challenges of developing visualization systems involving such text explanations and finding the right balance between these two approaches are still open [SBE* 18]. Here, we foresee an open research challenge upon how to combine visualizations, verbalization (text explanations), and voice commands (AI assistants) that should together perform overlapping tasks in complex visualization systems and propose task solutions to the users. As can be seen with our categoriza-

tion, analysts usually deal with data manipulation problems which can lead to compromising the trust, for example, (1) comparison of structures, (2) guidance in data selection, (3) outlier detection, (4) comparison of algorithms, and (5) in-situ comparison of concrete model structures. The aforementioned methods might provide a possible remedy for such compromises of trust.

Security vulnerabilities. When research is conducted in ML, there is always a factor that is not often taken into account at first: “how do we secure ML models from unethical attacks?” An instance of this idea is published by Ma et al. [MXLM20], explaining how visualization can assist in avoiding vulnerabilities of adversarial attacks in ML. Specifically, their focus is on how to avoid data poisoning attacks from the models, data instances, features, and local structures perspectives with the use of their VA approach. Nowadays, visualization systems are deployed online for users to access them easily. Such internet-accessibility leads to further problems concerning security vulnerabilities. This is one of the advantages of TensorFlow.js, which utilizes the WebGL-accelerated implementation of JavaScript in web browsers to implement and use ML models on local computers.

Fairness of the decisions. Going beyond interpretability towards more explainability is another open challenge. However, general proposals of frameworks in the visualization community combining ML and visualizations have been already described in recent research papers [MXQR19, SSSEA20]. These global frameworks were divided into smaller parts by other works that compare DL methods, for instance [MMD*19]. Further tools explore local trends instead of global patterns [ZWRH14]. Two further open questions reaching beyond interpretability and explainability are “how fair were all those decisions and what if we have chosen another path?” and “how can fairness be translated between the trust levels?” (cf. the work by Ahn and Lin [AL20]).

Ways of communication and collaboration. Increasing the users’ trust in ML models is not a trivial task. Visualization can assist in this challenge in multiple ways. A good starting point is employing simple techniques, such as querying specific data instances and areas of interest, in a user-friendly way [HNH*12]. However, the issue of improving trustworthiness in ML with visualization is also related to the issue of improving the trust for visualization itself [BRBF14, BBG19]. To achieve the best outcome when evaluating visualization designs, the input data, the goals, and the target group of a visualization should be under the spotlight. On the optimistic side, many papers exist that try to tackle the challenges of evaluation and design choices for visualizations [FAAM16, KPHL16, Kos16, LTBS*18, MSSW16, QH16]. Development of further guidelines and best practices for (1) how people within different scientific fields and varying backgrounds and experiences should communicate, and (2) which visualization techniques and systems should be established as a standardized interaction medium between them, present another open challenge. As previously discussed, Jentner et al. [JSS*18] suggest that metaphorical narratives can explain the ML models to various target groups in a user-friendly way, but further research is required in this regard.

Almost unexplored areas. Related to the non-trust level classes (which implicitly influence trust), we believe that all underrepresented categories can pose as new ideas for novel research. For ex-

ample, visualization researchers have still not provided sufficient support for some specific NNs, such as convolutional deep belief networks (CDBNs), deep residual networks (DRNs), and multi-column DNNs (MCDNNs). Also in ensemble learning, visualization tools that target solely the boosting techniques are quite rare, e.g., gradient boosting and adaptive boosting (AdaBoost) appear not to be covered to the same level as random forests. Another example of such a category is stacking ensemble learning, i.e., constructing a combination (a stack) of different models that should become the input for other meta-model(s). Employing visualization to facilitate the experts in developing and using such stacks in a trustworthy way without resorting to trial and error is also an open research challenge. Additionally, regression problems are also far less covered than classification. In unsupervised learning, association/pattern mining is uncommonly investigated by visual tools. To conclude this paragraph, reinforcement learning approaches are almost ignored with only a few available papers covering this area of how visualization can help to monitor an automatically controlled learning process [SPBA19]. In reinforcement learning, classification tasks, i.e., letting an agent act on the inputs and learn value functions [WvPS11], are not once addressed by visualization.

9. Conclusion

In this survey, we study the state of the art in enhancing trust in machine learning (ML) models with the use of visualizations. We introduced the background necessary for defining trustworthiness of ML models and explained the methodology used to select relevant papers in the literature. Based on the selected 200 peer-reviewed publications that introduce a large variety of visualization techniques to increase trust in ML models and their results, we proposed a fine-grained categorization comprising 8 high-level aspects partitioned into 18 category groups that on their part contain 119 categories in total. In addition, we performed a topic analysis to be able to discover connections and emerging topics among the 200 papers. Further analyses of the categorized data involved category correlations, temporal trends, and data sets used in the respective publications. In order to make our categorization and the assignment of papers into categories accessible for the public, an interactive survey browser—called TrustMLVis Browser—was implemented and made available online. It supports the readers of this STAR in the exploration of the rich information provided in this work, thus facilitating future research in enhancing trustworthiness of ML models with the help of interactive visualizations. Our findings indicate the growing interest for developing visualizations in ML to improve trustworthiness in the context of various data domains, tasks, and multidisciplinary applications. As future work, we intend to continue extending and refining the survey data set, categorization, and corresponding analyses, as well as maintaining the online survey browser.

References

- [AA97] ALIMOĞLU F., ALPAYDIN E.: Combining multiple representations and classifiers for pen-based handwritten digit recognition. In *Proceedings of the Fourth International Conference on Document Analysis and Recognition* (1997), vol. 2 of *ICDAR '97*, IEEE, pp. 637–640. doi:10.1109/ICDAR.1997.620583. 25

- [AASB19] ABBAS M. M., AUPETIT M., SEDLMAIR M., BENSMAIL H.: ClustMe: A visual quality measure for ranking monochrome scatterplots based on cluster patterns. *Computer Graphics Forum* 38, 3 (June 2019), 225–236. doi:10.1111/cgf.13684. 10
- [ACD*15] AMERSHI S., CHICKERING M., DRUCKER S. M., LEE B., SIMARD P., SUH J.: ModelTracker: Redesigning performance analysis tools for machine learning. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), CHI '15, ACM, pp. 337–346. doi:10.1145/2702123.2702509. 17
- [ACKK14] AMERSHI S., CAKMAK M., KNOX W. B., KULESZA T.: Power to the people: The role of humans in interactive machine learning. *AI Magazine* 35, 4 (Dec. 2014), 105–120. doi:10.1609/aimag.v35i4.2513. 8
- [AEM11] ALBUQUERQUE G., EISEMANN M., MAGNOR M.: Perception-based visual quality measures. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2011), VAST '11, IEEE, pp. 13–20. doi:10.1109/VAST.2011.6102437. 22
- [AGW18] ADEL T., GHAHRAMANI Z., WELLER A.: Discovering interpretable representations for both deep generative and discriminative models. In *Proceedings of the 35th International Conference on Machine Learning* (2018), vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 50–59. URL: <http://proceedings.mlr.press/v80/adell18a.html>. 7
- [AHH*14] ALSALLAKH B., HANBURY A., HAUSER H., MIKSCH S., RAUBER A.: Visual methods for analyzing probabilistic classification data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 1703–1712. doi:10.1109/TVCG.2014.2346660. 20
- [AJY*18] ALSALLAKH B., JOURABLOO A., YE M., LIU X., REN L.: Do convolutional neural networks learn class hierarchy? *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 152–162. doi:10.1109/TVCG.2017.2744683. 22
- [AK98] ALPAYDIN E., KAYNAK C.: Cascaded classifiers. *Kybernetika* 34, 4 (July 1998), 369–374. URL: <https://dml.cz/handle/10338.dmlcz/135217>. 25
- [AL20] AHN Y., LIN Y.: FairSight: Visual analytics for fairness in decision making. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 1086–1095. doi:10.1109/TVCG.2019.2934262. 17, 18, 28
- [AMJ18] ALVAREZ-MELIS D., JAAKKOLA T. S.: On the robustness of interpretability methods. In *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning* (2018), WHI '18. arXiv:1806.08049. 15
- [And36] ANDERSON E.: The species problem in Iris. *Annals of the Missouri Botanical Garden* 23, 3 (Sept. 1936), 457–509. URL: <http://jstor.org/stable/2394164>. 25, 26
- [Arr05] MIT-BIH Arrhythmia Database, 2005. Accessed January 10, 2020. URL: <https://sdo.gsfc.nasa.gov/>. 25
- [Art18] ARTICLE 29 DATA PROTECTION WORKING PARTY: Guidelines on automated individual decision-making and profiling for the purposes of Regulation 2016/679 (WP251rev.01), Feb. 2018. Accessed January 10, 2020. URL: https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053. 1
- [ASDG*00] ARGENZIANO G., SOYER H. P., DE GIORGIO V., PICCOLO D., CARLI P., DELFINO M., FERRARI A., HOFMANN-WELLENHOF R., MASSI D., MAZZOCCHETTI G., SCALVENZI M., WOLF I. H.: *Interactive Atlas of Dermoscopy*. Edra Medical Publishing and New Media, Milan, Italy, 2000. URL: <https://espace.library.uq.edu.au/view/UQ:229410>. 25
- [ASW*19] ARENDT D., SALDANHA E., WESSLEN R., VOLKOVA S., DOU W.: Towards rapid interactive machine learning: Evaluating trade-offs of classification without representation. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (2019), IUI '19, ACM, pp. 591–602. doi:10.1145/3301275.3302280. 17
- [Aup07] AUPETIT M.: Visualizing distortions and recovering topology in continuous projection techniques. *Neurocomputing* 70, 7–9 (Mar. 2007), 1304–1330. doi:10.1016/j.neucom.2006.11.018. 17, 22
- [AW12] AHMED Z., WEAVER C.: An adaptive parameter space-filling algorithm for highly interactive cluster exploration. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2012), VAST '12, IEEE, pp. 13–22. doi:10.1109/VAST.2012.6400493. 10
- [AYMW11] AHMED Z., YOST P., MCGOVERN A., WEAVER C.: Steerable clustering for visual analysis of ecosystems. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2011), EuroVA '11, The Eurographics Association. doi:10.2312/PE/EuroVAST/EuroVA11/049–052. 17
- [BAF*14] BÖGL M., AIGNER W., FILZMOSER P., GSCHWANDTNER T., LAMMARSCH T., MIKSCH S., RIND A.: Visual analytics methods to guide diagnostics for time series model predictions. In *Proceedings of the IEEE VIS Workshop on Visualization for Predictive Analytics* (2014), VPA '14. URL: <http://predictive-workshop.github.io/>. 19, 20
- [BAF*15] BÖGL M., AIGNER W., FILZMOSER P., GSCHWANDTNER T., LAMMARSCH T., MIKSCH S., RIND A.: Integrating predictions in time series model selection. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2015), EuroVA '15, The Eurographics Association. doi:10.2312/eurova.201511107. 20
- [BAL*15] BROOKS M., AMERSHI S., LEE B., DRUCKER S. M., KAPOOR A., SIMARD P.: FeatureInsight: Visual support for error-driven feature ideation in text classification. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2015), VAST '15, IEEE, pp. 105–112. doi:10.1109/VAST.2015.7347637. 21
- [BAPB*16] BOUDJELOUD-ASSALA L., PINHEIRO P., BLANSCHÉ A., TAMISIER T., OTJACQUES B.: Interactive and iterative visual clustering. *Information Visualization* 15, 3 (2016), 181–197. doi:10.1177/1473871615571951. 21
- [BBG19] BÖRNER K., BUECKLE A., GINDA M.: Data visualization literacy: Definitions, conceptual frameworks, exercises, and assessments. *Proceedings of the National Academy of Sciences* 116, 6 (2019), 1857–1864. doi:10.1073/pnas.1807180116. 7, 28
- [BDSF17] BERNARD J., DOBERMANN E., SEDLMAIR M., FELLNER D. W.: Combining cluster and outlier analysis with visual analytics. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2017), EuroVA '17, The Eurographics Association. doi:10.2312/eurova.20171114. 21
- [Bec16] BECKER K.: Identifying the gender of a voice using machine learning, 2016. Accessed January 10, 2020. URL: <http://primaryobjects.com/2016/06/22/identifying-the-gender-of-a-voice-using-machine-learning/>. 25
- [BEF17] BADAM S. K., ELMQVIST N., FEKETE J.-D.: Steering the craft: UI elements and visualizations for supporting progressive visual analytics. *Computer Graphics Forum* 36, 3 (June 2017), 491–502. doi:10.1111/cgf.13205. 21
- [Bel03] BELLMAN R. E.: *Dynamic Programming*. Dover Publications, Inc., Mineola, NY, USA, 2003. 27
- [Bes12] Best City Contest, 2012. Accessed January 10, 2020. URL: <http://eiu2012contest.blogspot.com/>. 25
- [BFOS84] BREIMAN L., FRIEDMAN J. H., OLSEN R. A., STONE C. J.: *Classification and Regression Trees*. The Wadsworth Statistics/Probability Series. Wadsworth & Brooks/Cole Advanced Books & Software, Monterey, CA, USA, 1984. URL: <https://cds.cern.ch/record/2253780>. 25
- [BHGK14] BEHAM M., HERZNER W., GRÖLLER M. E., KEHRER J.: Cupid: Cluster-based exploration of geometry generators with parallel coordinates and radial trees. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 1693–1702. doi:10.1109/TVCG.2014.2346626. 10

- [BHI09] BASTIAN M., HEYMANN S., JACOMY M.: Gephi: An open source software for exploring and manipulating networks. In *Proceedings of the International AAAI Conference on Weblogs and Social Media* (2009), ICWSM '09, AAAI, pp. 361–362. URL: <https://aaai.org/ocs/index.php/ICWSM/09/paper/view/154>. 11
- [BHK97] BELHUMEUR P. N., HESPANHA J. P., KRIEGMAN D. J.: Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 7 (July 1997), 711–720. doi:10.1109/34.598228. 25
- [BHR*19] BERNARD J., HUTTER M., RITTER C., LEHMANN M., SEDLMAIR M., ZEPPELZAUER M.: Visual analysis of degree-of-interest functions to support selection strategies for instance labeling. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2019), EuroVA '19, The Eurographics Association. doi:10.2312/eurova.20191116. 22
- [BHZ*18] BERNARD J., HUTTER M., ZEPPELZAUER M., FELLNER D., SEDLMAIR M.: Comparing visual-interactive labeling with active learning: An experimental study. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 298–308. doi:10.1109/TVCG.2017.2744818. 10
- [Bir06] BIRD S.: NLTK: The natural language toolkit. In *Proceedings of the COLING/ACL — Interactive Presentation Sessions* (2006), COLING-ACL '06, ACL, pp. 69–72. doi:10.3115/1225403.1225421. 19
- [BKSS14] BEHRISCH M., KORKMAZ F., SHAO L., SCHRECK T.: Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2014), VAST '14, IEEE, pp. 43–52. doi:10.1109/VAST.2014.7042480. 10
- [BKW16] BECK F., KOCH S., WEISKOPF D.: Visual analysis and dissemination of scientific literature collections with SurVis. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 180–189. doi:10.1109/TVCG.2015.2467757. 27
- [BL09] BERTINI E., LALANNE D.: Surveying the complementary role of automatic data analysis and visualization in knowledge discovery. In *Proceedings of the ACM SIGKDD Workshop on Visual Analytics and Knowledge Discovery: Integrating Automated Analysis with Interactive Exploration* (2009), VAKD '09, ACM, pp. 12–20. doi:10.1145/1562849.1562851. 2
- [BLBC12] BROWN E. T., LIU J., BRODLEY C. E., CHANG R.: Dis-Function: Learning distance functions interactively. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2012), VAST '12, IEEE, pp. 83–92. doi:10.1109/VAST.2012.6400486. 22
- [BM92] BENNETT K. P., MANGASARIAN O. L.: Robust linear programming discrimination of two linearly inseparable sets. *Optimization Methods and Software* 1, 1 (Apr. 1992), 23–34. doi:10.1080/10556789208805504. 25
- [BNJ03] BLEI D. M., NG A. Y., JORDAN M. I.: Latent Dirichlet allocation. *Journal of Machine Learning Research* 3 (Mar. 2003), 993–1022. URL: <http://jmlr.org/papers/v3/blei03a.html>. 19
- [BPD16] BPD Field Interrogation and Observation (FIO) dataset, 2016. Accessed January 10, 2020. URL: <https://data.boston.gov/dataset/boston-police-department-fio>. 25
- [BPF11] BERGER W., PIRINGER H., FILZMOSER P., GRÖLLER E.: Uncertainty-aware exploration of continuous parameter spaces using multivariate prediction. *Computer Graphics Forum* 30, 3 (June 2011), 911–920. doi:10.1111/j.1467-8659.2011.01940.x. 18
- [BR10] BROEMSTRUP T., REUTER N.: Molecular dynamics simulations of mixed acidic/zwitterionic phospholipid bilayers. *Biophysical Journal* 99, 3 (Aug. 2010), 825–833. doi:10.1016/j.bpj.2010.04.064. 25
- [BRBF14] BOY J., RENSINK R. A., BERTINI E., FEKETE J.-D.: A principled way of assessing visualization literacy. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 1963–1972. doi:10.1109/TVCG.2014.2346984. 7, 28
- [BTB14] BRUNI E., TRAN N. K., BARONI M.: Multimodal distributional semantics. *Journal of Artificial Intelligence Research* 49, 1 (Jan. 2014), 1–47. doi:10.1613/jair.4135. 25
- [BvLBS11] BREMM S., VON LANDESBERGER T., BERNARD J., SCHRECK T.: Assisted descriptor selection based on visual comparative data analysis. *Computer Graphics Forum* 30, 3 (June 2011), 891–900. doi:10.1111/j.1467-8659.2011.01938.x. 13, 15
- [BWZ*20] BORLAND D., WANG W., ZHANG J., SHRESTHA J., GOTZ D.: Selection bias tracking and detailed subset comparison for high-dimensional data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 429–439. doi:10.1109/TVCG.2019.2934209. 17
- [BZL*18] BERNARD J., ZEPPELZAUER M., LEHMANN M., MÜLLER M., SEDLMAIR M.: Towards user-centered active learning algorithms. *Computer Graphics Forum* 37, 3 (June 2018), 121–132. doi:10.1111/cgf.13406. 16
- [BZSA18] BERNARD J., ZEPPELZAUER M., SEDLMAIR M., AIGNER W.: VIAL: A unified process for visual interactive labeling. *The Visual Computer* 34, 9 (Sept. 2018), 1189–1207. doi:10.1007/s00371-018-1500-3. 10
- [CAA*19] CHEN S., ANDRIENKO N., ANDRIENKO G., ADILOVA L., BARLET J., KINDERMANN J., NGUYEN P. H., THONNARD O., TURKAY C.: LDA ensembles for interactive exploration and categorization of behaviors. *IEEE Transactions on Visualization and Computer Graphics* (2019). doi:10.1109/TVCG.2019.2904069. 20
- [Cad09] Cadaster Challenge, 2009. Accessed January 10, 2020. URL: <http://www.cadaster.eu/node/67.html>. 25
- [CBB*19] CHEGINI M., BERNARD J., BERGER P., SOURIN A., ANDREWS K., SCHRECK T.: Interactive labelling of a multivariate dataset for supervised machine learning using linked visualisations, clustering, and active learning. *Visual Informatics* 3, 1 (Mar. 2019), 9–17. Proceedings of PacificVAST 2019. doi:10.1016/j.visinf.2019.03.002. 17
- [CBK09] CHANDOLA V., BANERJEE A., KUMAR V.: Anomaly detection: A survey. *ACM Computing Surveys* 41, 3 (July 2009). doi:10.1145/1541880.1541882. 6
- [CBY10] CHEN Y., BARLOWE S., YANG J.: Click2Annotate: Automated insight externalization with rich semantics. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2010), VAST '10, IEEE, pp. 155–162. doi:10.1109/VAST.2010.5652885. 11
- [CCA*09] CORTEZ P., CERDEIRA A., ALMEIDA F., MATOS T., REIS J.: Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems* 47, 4 (Nov. 2009), 547–553. Smart Business Networks: Concepts and Empirical Evidence. doi:10.1016/j.dss.2009.05.016. 25, 26
- [CCR*19] CHOI I. K., CHILDERS T., RAVEENDRANATH N. K., MISHRA S., HARRIS K., REDA K.: Concept-driven visual analytics: An exploratory study of model- and hypothesis-based reasoning with visualizations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2019), CHI '19, ACM, pp. 68:1–68:14. doi:10.1145/3290605.3300298. 11
- [CCZ*16] CHEN Y., CHEN Q., ZHAO M., BOYER S., VEERAMACHANENI K., QU H.: DropoutSeer: Visualizing learning patterns in massive open online courses for dropout reasoning and prediction. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2016), VAST '16, IEEE, pp. 111–120. doi:10.1109/VAST.2016.7883517. 20
- [CD92] CHASE M. A., DUMMER G. M.: The role of sports as a social status determinant for children. *Research Quarterly for Exercise and Sport* 63, 4 (Dec. 1992), 418–424. doi:10.1080/02701367.1992.10608764. 25
- [CD18a] CAVALLO M., DEMIRALP Ç.: Track Xplorer: A system for visual analysis of sensor-based motor activity predictions. *Computer Graphics Forum* 37, 3 (June 2018), 339–349. doi:10.1111/cgf.13424. 17

- [CD18b] CAVALLO M., DEMIRALP Ç.: A visual interaction framework for dimensionality reduction based data exploration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018), CHI '18, ACM. doi:10.1145/3173574.3174209. 20
- [CD19] CAVALLO M., DEMIRALP Ç.: Clustrophile 2: Guided visual clustering analysis. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 267–276. doi:10.1109/TVCG.2018.2864477. 19
- [CDF*98] CRAVEN M., DIPASQUO D., FREITAG D., MCCALLUM A., MITCHELL T., NIGAM K., SLATTERY S.: Learning to extract symbolic knowledge from the World Wide Web. In *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence* (1998), AAAI '98/IAAI '98, American Association for Artificial Intelligence, pp. 509–516. doi:10.5555/295240.295725. 25
- [CDPP*17] CRESCI S., DI PIETRO R., PETROCCHI M., SPOGNARDI A., TESCONI M.: The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *Proceedings of the 26th International Conference on World Wide Web Companion* (2017), WWW '17 Companion, International World Wide Web Conferences Steering Committee, pp. 963–972. doi:10.1145/3041021.3055135. 25
- [CDS09] CROSSNO P. J., DUNLAVY D. M., SHEAD T. M.: LSAView: A tool for visual exploration of latent semantic modeling. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2009), VAST '09, IEEE, pp. 83–90. doi:10.1109/VAST.2009.5333428. 17, 21
- [CEH*19] CABRERA Á. A., EPPERSON W., HOHMAN F., KAHNG M., MORGENSTERN J., CHAU D. H.: FairVis: Visual analytics for discovering intersectional bias in machine learning. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2019), VAST '19, IEEE. arXiv:1904.05419. 18
- [CGF12] CETTOLO M., GIRARDI C., FEDERICO M.: WIT³: Web inventory of transcribed and translated talks. In *Proceedings of the 16th Annual Conference of the European Association for Machine Translation* (May 2012), EAMT '12, EAMT, pp. 261–268. URL: <http://mt-archive.info/EAMT-2012-complete.pdf>. 25
- [CGM19] CENEDA D., GSCHWANDTNER T., MIKSCH S.: A review of guidance approaches in visual data analysis: A multifocal perspective. *Computer Graphics Forum* 38, 3 (2019), 861–879. doi:10.1111/cgf.13730. 6
- [CGR*17] CHAE J., GAO S., RAMANTHAN A., STEED C., TOURASSI G. D.: Visualization for classification in deep neural networks. In *Proceedings of the Workshop on Visual Analytics for Deep Learning* (2017), VADL '17. URL: <https://vadl2017.github.io/>. 16
- [CHAS18] CUTURA R., HOLZER S., AUPETIT M., SEDLMAIR M.: VisCoDeR: A tool for visually comparing dimensionality reduction algorithms. In *Proceedings of the European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* (Jan. 2018), ESANN '18, Ciaco - i6doc.com, pp. 105–110. URL: <https://www.elen.ucl.ac.be/esann/proceedings/papers.php?ann=2018.17>
- [CHH*19] CASHMAN D., HUMAYOUN S. R., HEIMERL F., PARK K., DAS S., THOMPSON J. R., SAKET B., MOSCA A., STASKO J., ENDERT A., GLEICHER M., CHANG R.: A user-based visual analytics workflow for exploratory model analysis. *Computer Graphics Forum* 38, 3 (June 2019), 185–199. doi:10.1111/cgf.13681. 17
- [CHPY06] COHEN A. M., HERSH W. R., PETERSON K., YEN P.-Y.: Reducing workload in systematic review preparation using automated citation classification. *Journal of the American Medical Informatics Association* 13, 2 (Mar. 2006), 206–219. doi:10.1197/jamia.M1929. 25
- [CJH19] CAI C. J., JONGEJAN J., HOLBROOK J.: The effects of example-based explanations in a machine learning interface. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (2019), IUI '19, ACM, pp. 258–262. doi:10.1145/3301275.3302289. 2
- [CL18] CHOO J., LIU S.: Visual analytics for explainable deep learning. *IEEE Computer Graphics and Applications* 38, 4 (July 2018), 84–92. doi:10.1109/MCG.2018.042731661. 9, 13
- [CLKP10] CHOO J., LEE H., KIHM J., PARK H.: iVisClassifier: An interactive visual analytics system for classification based on supervised dimension reduction. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2010), VAST '10, IEEE, pp. 27–34. doi:10.1109/VAST.2010.5652443. 17
- [CM07] CORTEZ P., MORAIS A.: A data mining approach to predict forest fires using meteorological data. In *New Trends in Artificial Intelligence: Proceedings of the 13th Portuguese Conference on Artificial Intelligence* (2007), EPIA '07, APPIA, pp. 512–523. 25
- [CMJK20] CHATZIMPARMAS A., MARTINS R. M., JUSUFI I., KERREN A.: A survey of surveys on the use of visualization for interpreting machine learning models. *Information Visualization* (2020). doi:10.1177/1473871620904671. 8
- [CMN*16] COIMBRA D. B., MARTINS R. M., NEVES T. T., TELEA A. C., PAULOVICH F. V.: Explaining three-dimensional dimensionality reduction plots. *Information Visualization* 15, 2 (Apr. 2016), 154–172. doi:10.1177/1473871615600010. 16, 22
- [Coc11] COCK D. D.: Ames, Iowa: Alternative to the Boston Housing Data as an end of semester regression project. *Journal of Statistics Education* 19, 3 (Nov. 2011). doi:10.1080/10691898.2011.11889627. 25
- [Coh88] COHEN L. H.: Measurement of life events. In *Life Events and Psychological Functioning: Theoretical and Methodological Issues*. SAGE Publications, Thousand Oaks, CA, USA, 1988, pp. 11–30. 22
- [COM19] COMPAS recidivism risk score data and analysis—ProPublica, 2019. Accessed January 10, 2020. URL: <https://propublica.org/datastore/dataset/compas-recidivism-risk-score-data-and-analysis>. 25
- [CPCS20] CASHMAN D., PERER A., CHANG R., STROBELT H.: Abate, variate, and contemplate: Visual analytics for discovering neural architectures. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 863–873. doi:10.1109/TVCG.2019.2934261. 18
- [CPM*18] CASHMAN D., PATTERSON G., MOSCA A., WATTS N., ROBINSON S., CHANG R.: RNNbow: Visualizing learning via backpropagation gradients in RNNs. *IEEE Computer Graphics and Applications* 38, 6 (Nov. 2018), 39–50. doi:10.1109/MCG.2018.2878902. 15, 22
- [CRMH12] CHUANG J., RAMAGE D., MANNING C., HEER J.: Interpretation and trust: Designing model-driven visualizations for text analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2012), CHI '12, ACM, pp. 443–452. doi:10.1145/2207676.2207738. 3, 7
- [CRRS08] CHANG M.-W., RATINOV L.-A., ROTH D., SRIKUMAR V.: Importance of semantic representation: Dataless classification. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence* (2008), AAAI '08, AAAI Press. URL: <https://aaai.org/Library/AAAI/2008/aaai08-132.php>. 25
- [CS14] CHUANG J., SOCHER R.: Interactive visualizations for deep learning. In *Proceedings of the IEEE VIS Workshop on Visualization for Predictive Analytics* (2014), VPA '14. URL: <http://predictive-workshop.github.io/>. 17
- [CSG*18] CHEGINI M., SHAO L., GREGOR R., LEHMANN D. J., ANDREWS K., SCHRECK T.: Interactive visual exploration of local patterns in large scatterplot spaces. *Computer Graphics Forum* 37, 3 (June 2018), 99–109. doi:10.1111/cgf.13404. 17
- [CSV*18] CHEN N.-C., SUH J., VERWEY J., RAMOS G., DRUCKER S., SIMARD P.: AnchorViz: Facilitating classifier error discovery through interactive semantic data exploration. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces* (2018), IUI '18, ACM, pp. 269–280. doi:10.1145/3172944.3172950. 18

- [CSWJ18] CHEN J., SONG L., WAINWRIGHT M., JORDAN M.: Learning to explain: An information-theoretic perspective on model interpretation. In *Proceedings of the 35th International Conference on Machine Learning* (2018), vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 883–892. URL: <http://proceedings.mlr.press/v80/chen18j.html>. 7
- [CWGW19] CABALLERO H. S. G., WESTENBERG M. A., GEBRE B., WIJK J. J. V.: V-Awake: A visual analytics approach for correcting sleep predictions from deep learning models. *Computer Graphics Forum* 38, 3 (June 2019), 1–12. doi:10.1111/cgf.13667. 16
- [CWS*17] CYPKO M., WOJDAK J., STOEHR M., KIRCHNER B., PREIM B., DIETZ A., LEMKE H. U., OELTZE-JAFRA S.: Visual verification of cancer staging for therapy decision support. *Computer Graphics Forum* 36, 3 (June 2017), 109–120. doi:10.1111/cgf.13172. 20
- [Dal20] Descriptive machine Learning EXplanations (DALEX), 2020. Accessed January 10, 2020. URL: <https://modeloriented.github.io/DALEX/>. 2
- [Dar20] Defense Advanced Research Projects Agency — Explainable Artificial Intelligence (XAI) program information, 2020. Accessed January 10, 2020. URL: <https://darpa.mil/program/explainable-artificial-intelligence>. 1
- [dBD*12] DOS SANTOS AMORIM E. P., BRAZIL E. V., DANIELS J., JOIA P., NONATO L. G., SOUSA M. C.: iLAMP: Exploring high-dimensional spacing through backward multidimensional projection. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2012), VAST '12, IEEE, pp. 53–62. doi:10.1109/VAST.2012.6400489. 16
- [DCCE19] DAS S., CASHMAN D., CHANG R., ENDERT A.: BEAMES: Interactive multi-model steering, selection, and inspection for regression tasks. *IEEE Computer Graphics and Applications* 39, 9 (Sept. 2019). doi:10.1109/MCG.2019.2922592. 17
- [DDS*09] DENG J., DONG W., SOCHER R., LI L.-J., LI K., FEI-FEI L.: ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2009), CVPR '09, IEEE, pp. 248–255. doi:10.1109/CVPR.2009.5206848. 25, 26
- [DFP*20] DIMARA E., FRANCONERI S., PLAISANT C., BEZERIANOS A., DRAGICEVIC P.: A task-based taxonomy of cognitive biases for information visualization. *IEEE Transactions on Visualization and Computer Graphics* 26, 2 (Feb. 2020), 1413–1432. doi:10.1109/TVCG.2018.2872577. 7
- [DG17] DUA D., GRAFF C.: UCI Machine Learning Repository, 2017. URL: <http://archive.ics.uci.edu/ml>. 4, 25
- [DGL89] DUFF I. S., GRIMES R. G., LEWIS J. G.: Sparse matrix test problems. *ACM Transactions on Mathematical Software* 15, 1 (Mar. 1989), 1–14. doi:10.1145/62038.62043. 25
- [DK18] DUDLEY J. J., KRISTENSSON P. O.: A review of user interface design for interactive machine learning. *ACM Transactions on Interactive Intelligent Systems* 8, 2 (June 2018), 8:1–8:37. doi:10.1145/3185517. 8, 13
- [DLH20] DU M., LIU N., HU X.: Techniques for interpretable machine learning. *Communications of the ACM* 63, 1 (Jan. 2020), 68–77. doi:10.1145/3359786. 8
- [DP98] DIEKMANN R., PREIS R.: AG-Monien Graph, 1998. Accessed January 10, 2020. URL: http://cise.ufl.edu/research/sparse/matrices/AG-Monien/airfoill_dual.html. 25
- [EASKC18] EL-ASSADY M., SEVASTYANOVA R., KEIM D., COLLINS C.: ThreadReconstructor: Modeling reply-chains to untangle conversational text through visual analytics. *Computer Graphics Forum* 37, 3 (June 2018), 351–365. doi:10.1111/cgf.13425. 17
- [EC16] EUROPEAN PARLIAMENT, COUNCIL OF THE EUROPEAN UNION: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), Apr. 2016. Accessed January 10, 2020. URL: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>. 1
- [EDF87] EIN-DOR P., FELDMESSER J.: Attributes of the performance of central processing units: A relative performance prediction model. *Communications of the ACM* 30, 4 (Apr. 1987), 308–317. doi:10.1145/32232.32234. 25
- [EDF08] ELMQVIST N., DRAGICEVIC P., FEKETE J.-D.: Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (Nov. 2008), 1539–1148. doi:10.1109/TVCG.2008.153. 25
- [EGG*12] ENGEL D., GREFF K., GARTH C., BEIN K., WEXLER A., HAMANN B., HAGEN H.: Visual steering and verification of mass spectrometry data factorization in air quality research. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2275–2284. doi:10.1109/TVCG.2012.280. 19
- [EK09] EVANS A. M., KRUEGER J. I.: The psychology (and economics) of trust. *Social and Personality Psychology Compass* 3, 6 (Dec. 2009), 1003–1017. doi:10.1111/j.1751-9004.2009.00232.x. 3
- [ELSG08] ESS A., LEIBE B., SCHINDLER K., GOOL L. V.: A mobile vision system for robust multi-person tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2008), CVPR '08, IEEE. doi:10.1109/CVPR.2008.4587581. 25
- [Eng18] English Premier League players dataset, 2017/18, 2018. Accessed January 10, 2020. URL: <https://kaggle.com/mauryashubham/english-premier-league-players-dataset>. 25
- [ERT*17] ENDERT A., RIBARSKY W., TURKAY C., WONG B. W., NABNEY I., BLANCO I. D., ROSSI F.: The state of the art in integrating machine learning into visual analytics. *Computer Graphics Forum* 36, 8 (2017), 458–486. doi:10.1111/cgf.13092. 8
- [ESS18] European Social Survey (ESS), 2018. Accessed January 10, 2020. URL: <https://europeansocialsurvey.org/>. 25
- [Eva96] EVANS J. D.: *Straightforward Statistics for the Behavioral Sciences*. Brooks/Cole Publishing, Pacific Grove, CA, USA, 1996. 22
- [FAAM16] FEDERICO P., AMOR-AMORÓS A., MIKSCH S.: A nested workflow model for visual analytics design and validation. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), BELIV '16, ACM, pp. 104–111. doi:10.1145/2993901.2993915. 28
- [FBG19] FENG S., BOYD-GRABER J.: What can AI do for me?: Evaluating machine learning interpretations in cooperative play. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (2019), IUI '19, ACM, pp. 229–239. doi:10.1145/3301275.3302265. 2
- [FBT*10] FERDOSI B. J., BUDELMEIJER H., TRAGER S., WILKINSON M. H. F., ROERDINK J. B. T. M.: Finding and visualizing relevant subspaces for clustering high-dimensional astronomical data using connected morphological operators. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2010), VAST '10, IEEE, pp. 35–42. doi:10.1109/VAST.2010.5652450. 10
- [FBVV09] FREIRE A. L., BARRETO G. A., VELOSO M., VARELA A. T.: Short-term memory mechanisms in neural network learning of robot navigation tasks: A case study. In *Proceedings of the 6th Latin American Robotics Symposium* (2009), LARS '09, IEEE. doi:10.1109/LARS.2009.5418323. 25
- [FCS*20] FUJIWARA T., CHOU J., SHILPIKA S., XU P., REN L., MA K.-L.: An incremental dimensionality reduction method for visualizing streaming multidimensional data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 418–428. doi:10.1109/TVCG.2019.2934433. 19
- [FFFP04] FEI-FEI L., FERGUS R., PERONA P.: Learning generative visual models from few training examples: An incremental Bayesian

- approach tested on 101 object categories. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops* (2004), CVPRW '04, IEEE. doi:10.1109/CVPR.2004.383. 25
- [FGM*01] FINKELSTEIN L., GABRILOVICH E., MATIAS Y., RIVLIN E., SOLAN Z., WOLFMAN G., RUPPIN E.: Placing search in context: The concept revisited. In *Proceedings of the 10th International Conference on World Wide Web* (2001), WWW '01, ACM, pp. 406–414. doi:10.1145/371920.372094. 25
- [FKM20] FUJIWARA T., KWON O., MA K.-L.: Supporting analysis of dimensionality reduction results with contrastive learning. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 45–55. doi:10.1109/TVCG.2019.2934251. 19
- [FMH16] FRÖHLER B., MÖLLER T., HEINZL C.: GEMSe: Visualization-guided exploration of multi-channel segmentation algorithms. *Computer Graphics Forum* 35, 3 (June 2016), 191–200. doi:10.1111/cgf.12895. 19
- [Fri02] FRIEDMAN J. H.: Stochastic gradient boosting. *Computational Statistics & Data Analysis* 38, 4 (Feb. 2002), 367–378. doi:10.1016/S0167-9473(01)00065-2. 25
- [FSJ13] FERNSTAD S. J., SHAW J., JOHANSSON J.: Quality-based guidance for exploratory dimensionality reduction. *Information Visualization* 12, 1 (Jan. 2013), 44–64. doi:10.1177/1473871612460526. 15, 21
- [FT99] FOGG B. J., TSENG H.: The elements of computer credibility. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (1999), CHI '99, ACM, pp. 80–87. doi:10.1145/302979.303001. 3
- [FTG14] FANAEE-T H., GAMA J.: Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence* 2, 2 (June 2014), 113–127. doi:10.1007/s13748-013-0040-3. 25
- [FV14] FRENAY B., VERLEYSSEN M.: Classification in the presence of label noise: A survey. *IEEE Transactions on Neural Networks and Learning Systems* 25, 5 (May 2014), 845–869. doi:10.1109/TNNLS.2013.2292894. 6
- [FVC15] FERNANDES K., VINAGRE P., CORTEZ P.: A proactive intelligent decision support system for predicting the popularity of online news. In *Progress in Artificial Intelligence: Proceedings of the 17th Portuguese Conference on Artificial Intelligence (EPIA '15)* (2015), vol. 9273 of LNCS, Springer International Publishing, pp. 535–546. doi:10.1007/978-3-319-23485-4_53. 25
- [GBY*18] GILPIN L. H., BAU D., YUAN B. Z., BAJWA A., SPECTER M., KAGAL L.: Explaining explanations: An overview of interpretability of machine learning. In *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics* (2018), DSAA '18, IEEE, pp. 80–89. doi:10.1109/DSAA.2018.00018. 2
- [GC05] GREENE D., CUNNINGHAM P.: Producing accurate interpretable clusters from high-dimensional data. In *Knowledge Discovery in Databases: Proceedings of the 9th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD '05)* (2005), vol. 3721 of LNCS, Springer Berlin Heidelberg, pp. 486–494. doi:10.1007/11564126_49. 25
- [GDG11] GRGIC M., DELAC K., GRGIC S.: SCface — Surveillance Cameras Face Database. *Multimedia Tools Applications* 51, 3 (Feb. 2011), 863–879. doi:10.1007/s11042-009-0417-2. 25
- [GDM*19] GUO S., DU F., MALIK S., KOH E., KIM S., LIU Z., KIM D., ZHA H., CAO N.: Visualizing uncertainty and alternatives in event sequence predictions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), CHI '19, ACM, pp. 573:1–573:12. doi:10.1145/3290605.3300803. 22
- [GHG*19] GIL Y., HONAKER J., GUPTA S., MA Y., D'ORAZIO V., GARIJO D., GADEWAR S., YANG Q., JAHANSHAD N.: Towards human-guided machine learning. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (2019), IUI '19, ACM, pp. 614–624. doi:10.1145/3301275.3302324. 16, 21
- [GHP07] GRIFFIN G., HOLUB A., PERONA P.: Caltech-256 Object Category Dataset, 2007. URL: <https://resolver.caltech.edu/CaltechAUTHORS:CNS-TR-2007-001>. 25
- [GKN05] GANSNER E. R., KOREN Y., NORTH S. C.: Topological fish-eye views for visualizing large graphs. *IEEE Transactions on Visualization and Computer Graphics* 11, 4 (July 2005), 457–468. doi:10.1109/TVCG.2005.66. 25
- [Gle13] GLEICHER M.: Explainers: Expert explorations with crafted projections. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec. 2013), 2042–2051. doi:10.1109/TVCG.2013.157. 15
- [GM04] GABRILOVICH E., MARKOVITCH S.: Text categorization with many redundant features: Using aggressive feature selection to make SVMs competitive with C4.5. In *Proceedings of the 21st International Conference on Machine Learning* (2004), ICML '04, ACM. doi:10.1145/1015330.1015388. 25
- [GMP18] GOODFELLOW I., MCDANIEL P., PAPERNOT N.: Making machine learning robust against adversarial inputs. *Communications of the ACM* 61, 7 (June 2018), 56–66. doi:10.1145/3134599. 6
- [GNRM08] GARG S., NAM J. E., RAMAKRISHNAN I. V., MUELLER K.: Model-driven visual analytics. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2008), VAST '08, IEEE, pp. 19–26. doi:10.1109/VAST.2008.4677352. 19
- [Goo20] Google Cloud Explainable AI, 2020. Accessed January 10, 2020. URL: <https://cloud.google.com/explainable-ai/>. 2
- [GRM10] GARG S., RAMAKRISHNAN I. V., MUELLER K.: A visual analytics approach to model learning. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2010), VAST '10, IEEE, pp. 67–74. doi:10.1109/VAST.2010.5652484. 10
- [GRNT16] GRÜN F., RUPPRECHT C., NAVAB N., TOMBARI F.: A taxonomy and library for visualizing learned features in convolutional neural networks. In *Proceedings of the ICML Workshop on Visualization for Deep Learning* (2016), DL '16. arXiv:1606.07757. 8
- [GS04] GRIFFITHS T. L., STEYVERS M.: Finding scientific topics. *Proceedings of the National Academy of Sciences* 101, suppl 1 (2004), 5228–5235. doi:10.1073/pnas.0307752101. 19
- [GS14] GOTZ D., SUN J.: Visualizing accuracy to improve predictive model performance. *Proceedings of the IEEE VIS Workshop on Visualization for Predictive Analytics* (2014). URL: <http://predictive-workshop.github.io/>. 21
- [GSC16] GOTZ D., SUN S., CAO N.: Adaptive contextualization: Combating bias during high-dimensional visualization and data selection. In *Proceedings of the 21st International Conference on Intelligent User Interfaces* (2016), IUI '16, ACM, pp. 85–95. doi:10.1145/2856767.2856779. 21, 27
- [GSK*20] GEHRMANN S., STROBELT H., KRÜGER R., PFISTER H., RUSH A. M.: Visual interaction with deep learning models through collaborative semantic inference. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 884–894. doi:10.1109/TVCG.2019.2934595. 17
- [GSS*20] GÖRTLER J., SPINNER T., STREEB D., WEISKOPF D., DEUSSEN O.: Uncertainty-aware principal component analysis. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 822–831. doi:10.1109/TVCG.2019.2934812. 17
- [GTdS*18] GARCIA R., TELEA A. C., DA SILVA B. C., TØRRESEN J., AO LUIZ DIHL COMBA J.: A task-and-technique centered survey on visual analytics for deep learning model engineering. *Computers & Graphics* 77 (2018), 30–49. doi:10.1016/j.cag.2018.09.018. 9
- [GTS*08] GOETZ C. G., TILLEY B. C., SHAFTMAN S. R., STEBBINS G. T., FAHN S., MARTINEZ-MARTIN P., POEWE W., SAMPAIO C., STERN M. B., DODEL R., DUBOIS B., HOLLOWAY R., JANKOVIC J., KULISEVSKY J., LANG A. E., LEES A., LEURGANS S., LEWITT P. A., NYENHUIS D., OLANOW C. W., RASCOL O., SCHRAG A.,

- TERESI J. A., VAN HILTEN J. J., LAPPELLE N.: Movement Disorder Society-sponsored revision of the Unified Parkinson's Disease Rating Scale (MDS-UPDRS): Scale presentation and clinimetric testing results. *Movement Disorders* 23, 15 (Nov. 2008), 2129–2170. doi:10.1002/mds.22340. 25
- [HB15] HOFF K. A., BASHIR M.: Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors* 57, 3 (May 2015), 407–434. doi:10.1177/0018720814547570. 3
- [HDK*19] HAMID S., DERSTROFF A., KLEMM S., NGO Q. Q., JIANG X., LINSEN L.: Visual ensemble analysis to study the influence of hyper-parameters on training deep neural networks. In *Proceedings of the EuroVis Workshop on Machine Learning Methods in Visualisation for Big Data* (2019), MLVis '19, The Eurographics Association. doi:10.2312/mlvis.20191160. 18, 22
- [HGC15] HIGUERA C., GARDINER K. J., CIOS K. J.: Self-organizing feature maps identify proteins critical to learning in a mouse model of Down syndrome. *PLOS ONE* 10, 6 (June 2015). doi:10.1371/journal.pone.0129126. 25
- [HHC*19] HOHMAN F., HEAD A., CARUANA R., DELINE R., DRUCKER S. M.: Gamut: A design probe to understand how data scientists understand machine learning models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), CHI '19, ACM, pp. 579:1–579:13. doi:10.1145/3290605.3300809. 18
- [Hic14] HICKEY W.: A statistical analysis of the work of Bob Ross, 2014. Accessed January 10, 2020. URL: <https://fivethirtyeight.com/features/a-statistical-analysis-of-the-work-of-bob-ross/>. 25
- [HJBU13] HOFFMAN R. R., JOHNSON M., BRADSHAW J. M., UNDERBRINK A.: Trust in automation. *IEEE Intelligent Systems* 28, 1 (Jan.–Feb. 2013), 84–88. doi:10.1109/MIS.2013.24. 3
- [HKPC19] HOHMAN F., KAHNG M., PIENTA R., CHAU D. H.: Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE Transactions on Visualization and Computer Graphics* 25, 8 (Aug. 2019), 2674–2693. doi:10.1109/TVCG.2018.2843369. 9, 13
- [HLW*20] HAZARIKA S., LI H., WANG K., SHEN H., CHOU C.: NNVA: Neural network assisted visual analysis of yeast cell polarization simulation. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 34–44. doi:10.1109/TVCG.2019.2934591. 17, 22
- [HNN*12] HÖFERLIN B., NETZEL R., HÖFERLIN M., WEISKOPF D., HEIDEMANN G.: Inter-active learning of ad-hoc classifiers for video visual analytics. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2012), VAST '12, IEEE, pp. 23–32. doi:10.1109/VAST.2012.6400492. 19, 28
- [How18] HOWARD A.: Investigations into the human-AI trust phenomenon. Plenary invited talk at NeurIPS '18, Dec. 2018. 4
- [HPRC20] HOHMAN F., PARK H., ROBINSON C., CHAU D. H.: Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 1096–1106. doi:10.1109/TVCG.2019.2934659. 22
- [HR78] HARRISON D., RUBINFELD D. L.: Hedonic housing prices and the demand for clean air. *Journal of Environmental Economics and Management* 5, 1 (Mar. 1978), 81–102. doi:10.1016/0095-0696(78)90006-2. 25
- [HRK15] HILL F., REICHAERT R., KORHONEN A.: SimLex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics* 41, 4 (Dec. 2015), 665–695. doi:10.1162/COLI_a_00237. 25
- [HSD19] HOHMAN F., SRINIVASAN A., DRUCKER S. M.: TeleGam: Combining visualization and verbalization for interpretable machine learning. In *2019 IEEE Visualization Conference (VIS)* (Oct 2019), pp. 151–155. doi:10.1109/VISUAL.2019.8933695. 15, 17
- [HSPC06] HAWKES E. R., SANKARAN R., PÉBAY P. P., CHEN J. H.: Direct numerical simulation of ignition front propagation in a constant volume with temperature inhomogeneities: II. Parametric study. *Combustion and Flame* 145, 1–2 (Apr. 2006), 145–159. doi:10.1016/j.combustflame.2005.09.018. 25
- [Huu00] HUUSKONEN J.: Estimation of aqueous solubility for a diverse set of organic compounds based on molecular topology. *Journal of Chemical Information and Computer Sciences* 40, 3 (May 2000), 773–777. doi:10.1021/ci9901338. 25
- [HV81] HENDERSON H. V., VELLEMAN P. F.: Building multiple regression models interactively. *Biometrics* 37, 2 (June 1981), 391–411. doi:10.2307/2530428. 25
- [HVP*19] HÖLLT T., VILANOVA A., PEZZOTTI N., LELIEVELDT B., HAUSER H.: Focus+context exploration of hierarchical embeddings. *Computer Graphics Forum* 38, 3 (June 2019), 569–579. doi:10.1111/cgf.13711. 15
- [Hyp19] Hyperspectral remote sensing scenes, 2019. Accessed January 10, 2020. URL: http://ehu.eus/ccwintco/index.php?title=Hyperspectral_Remote_Sensing_Scenes. 25
- [i-L19] i-Lids multicamera tracking — UK government, 2019. Accessed January 10, 2020. URL: <http://homeoffice.gov.uk/science-research/hosdb/i-lids/>. 25
- [Ima19] ImageCLEF — The CLEF cross language image retrieval track, 2019. Accessed January 10, 2020. URL: <https://imageclef.org/>. 25
- [IMI*10] INGRAM S., MUNZNER T., IRVINE V., TORY M., BERGNER S., MÖLLER T.: DimStiller: Workflows for dimensional analysis and reduction. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2010), VAST '10, IEEE, pp. 3–10. doi:10.1109/VAST.2010.5652392. 16
- [Inf17] InfoVis and VAST papers, 2017. Accessed January 10, 2020. URL: <https://cc.gatech.edu/gvu/ii/jigsaw/datafiles.html>. 25
- [JC17a] JASSBY A. D., CLOERN J. E.: WQ: Exploring water quality monitoring data, 2017. Accessed January 10, 2020. URL: <https://cran.rstudio.com/web/packages/wq/>. 25
- [JC17b] JIANG B., CANNY J.: Interactive machine learning via a GPU-accelerated toolkit. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces* (2017), IUI '17, ACM, pp. 535–546. doi:10.1145/3025171.3025172. 19
- [JHB*17] JÄCKLE D., HUND M., BEHRISCH M., KEIM D. A., SCHRECK T.: Pattern Trails: Visual analysis of pattern transitions in subspaces. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2017), VAST '17, IEEE, pp. 1–12. doi:10.1109/VAST.2017.8585613. 22
- [JJ09] JOHANSSON S., JOHANSSON J.: Interactive dimensionality reduction through user-defined combinations of quality metrics. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (Nov. 2009), 993–1000. doi:10.1109/TVCG.2009.153. 16
- [JKM12] JANKOWSKA M., KEŠELJ V., MILIOS E.: Relative N-gram signatures: Document visualization at the level of character N-grams. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2012), VAST '12, IEEE, pp. 103–112. doi:10.1109/VAST.2012.6400484. 13
- [JPN15] JOIA P., PETRONETTO F., NONATO L.: Uncovering representative groups in multidimensional projections. *Computer Graphics Forum* 34, 3 (June 2015), 281–290. doi:10.1111/cgf.12640. 22
- [JRK*16] JONGEJAN J., ROWLEY H., KAWASHIMA T., KIM J., FOXGIEG N.: Quick, Draw! by Google Creative Lab, 2016. Accessed January 10, 2020. URL: <https://experiments.withgoogle.com/quick-draw>. 25
- [JSO19] JANIK A., SANKARAN K., ORTIZ A.: Interpreting black-box

- semantic segmentation models in remote sensing applications. In *Proceedings of the EuroVis Workshop on Machine Learning Methods in Visualisation for Big Data* (2019), MLVis '19, The Eurographics Association. doi:10.2312/mlvis.20191158. 1, 19
- [JSR*19] JI X., SHEN H., RITTER A., MACHIRAJU R., YEN P.: Visual exploration of neural document embedding in information retrieval: Semantics and feature selection. *IEEE Transactions on Visualization and Computer Graphics* 25, 6 (June 2019), 2181–2192. doi:10.1109/TVCG.2019.2903946. 17
- [JSS*18] JENTNER W., SEVASTJANOVA R., STOFFEL F., KEIM D. A., BERNARD J., EL-ASSADY M.: Minions, sheep, and fruits: Metaphorical narratives to explain artificial intelligence and build trust. In *Proceedings of the IEEE VIS Workshop on Visualization for AI Explainability* (2018), VISxAI '18. URL: <https://visxai.io/>. 19, 28
- [JZF*09] JEONG D. H., ZIEMKIEWICZ C., FISHER B., RIBARSKY W., CHANG R.: iPCA: An interactive system for PCA-based visual analytics. *Computer Graphics Forum* 28, 3 (June 2009), 767–774. doi:10.1111/j.1467-8659.2009.01475.x. 15
- [KAKC18] KAHNG M., ANDREWS P. Y., KALRO A., CHAU D. H.: ActiVis: Visual exploration of industry-scale deep neural network models. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 88–97. doi:10.1109/TVCG.2017.2744718. 15
- [KBR84] KONONENKO I., BRATKO I., ROŠKAR E.: Experiments in automatic learning of medical diagnostic rules. In *Proceedings of the International School for the Synthesis of Expert Knowledge Workshop* (1984). 25
- [KBWS15] KULESZA T., BURNETT M., WONG W.-K., STUMPF S.: Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (2015), IUI '15, ACM, pp. 126–137. doi:10.1145/2678025.2701399. 13
- [KC19] KAHNG M., CHAU D. H.: How does visualization help people learn deep learning? Evaluation of GAN Lab. In *Proceedings of IEEE VIS Workshop on Evaluation of Interactive Visual Machine Learning Systems* (2019), EVIVA-ML '19. URL: <https://eviva-ml.github.io/>. 18
- [KCK17] KIM K., CARLIS J. V., KEEFE D. F.: Comparison techniques utilized in spatial 3D and 4D data visualizations: A survey and future directions. *Computers & Graphics* 67 (2017), 138–147. doi:10.1016/j.cag.2017.05.005. 6
- [KCK*19] KWON B. C., CHOI M., KIM J. T., CHOI E., KIM Y. B., KWON S., SUN J., CHOO J.: RetainVis: Visual analytics with interpretable and interactive recurrent neural networks on electronic medical records. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 299–309. doi:10.1109/TVCG.2018.2865027. 17
- [KDFB16] KRAUSE J., DASGUPTA A., FEKETE J.-D., BERTINI E.: SeekAView: An intelligent dimensionality reduction strategy for navigating high-dimensional data spaces. In *Proceedings of the IEEE Symposium on Large Data Analysis and Visualization* (2016), LDAH '16, IEEE, pp. 11–19. doi:10.1109/LDAH.2016.7874305. 21
- [KDS*17] KRAUSE J., DASGUPTA A., SWARTZ J., APHINYANAPHONGS Y., BERTINI E.: A workflow for visual diagnostics of binary classifiers using instance-level explanations. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2017), VAST '17, IEEE, pp. 162–172. doi:10.1109/VAST.2017.8585720. 1, 15
- [KEV*18] KWON B. C., EYSENBACH B., VERMA J., NG K., DE FILIPPI C., STEWART W. F., PERER A.: Clustervision: Visual supervision of unsupervised clustering. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 142–151. doi:10.1109/TVCG.2017.2745085. 21
- [KFC16] KAHNG M., FANG D., CHAU D. H.: Visual exploration of machine learning results using data cube analysis. In *Proceedings of the Workshop on Human-In-the-Loop Data Analytics* (2016), HILDA '16, ACM, pp. 1:1–1:6. doi:10.1145/2939502.2939503. 19
- [KHP*11] KANDEL S., HEER J., PLAISANT C., KENNEDY J., VAN HAM F., RICHE N. H., WEAVER C., LEE B., BRODBECK D., BUONO P.: Research directions in data wrangling: Visualizations and transformations for usable and credible data. *Information Visualization* 10, 4 (Oct. 2011), 271–288. doi:10.1177/1473871611415994. 6
- [KJR*18] KAUER T., JOGLEKAR S., REDY M., AIELLO L. M., QUERCIA D.: Mapping and visualizing deep-learning urban beautification. *IEEE Computer Graphics and Applications* 38, 5 (Sept. 2018), 70–83. doi:10.1109/MCG.2018.053491732. 19
- [KK14] KUCHER K., KERREN A.: Text visualization browser: A visual survey of text visualization techniques. In *Poster Abstracts of IEEE VIS* (2014). 27
- [KK15] KUCHER K., KERREN A.: Text visualization techniques: Taxonomy, visual survey, and community insights. In *Proceedings of the 8th IEEE Pacific Visualization Symposium* (2015), PacificVis '15, IEEE, pp. 117–121. doi:10.1109/PACIFICVIS.2015.7156366. 13, 27
- [KKB19] KINKELDEY C., KORJAKOW T., BENJAMIN J. J.: Towards supporting interpretability of clustering results with uncertainty visualization. In *Proceedings of the EuroVis Workshop on Trustworthy Visualization* (2019), TrustVis '19, The Eurographics Association. doi:10.2312/trvis.20191183. 16
- [KKK14] KIM L., KIM J.-A., KIM S.: A guide for the utilization of Health Insurance Review and Assessment Service National Patient Samples. *Epidemiology and Health* 36 (July 2014). doi:10.4178/epih/e2014008. 25
- [KKS*19] KELLY C. J., KARTHESALINGAM A., SULEYMAN M., CORRADO G., KING D.: Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine* 17, 1 (2019), 195. doi:10.1186/s12916-019-1426-2. 1
- [KKW*17] KWON B. C., KIM H., WALL E., CHOO J., PARK H., ENDERT A.: AxiSketcher: Interactive nonlinear axis mapping of visualizations through user drawings. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 221–230. doi:10.1109/TVCG.2016.2598446. 15
- [KKZE20] KHAYAT M., KARIMZADEH M., ZHAO J., EBERT D. S.: VASSL: A visual analytics toolkit for social spambot labeling. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 874–883. doi:10.1109/TVCG.2019.2934266. 16, 17, 20
- [KLTH10] KAPOOR A., LEE B., TAN D., HORVITZ E.: Interactive optimization for steering machine classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2010), CHI '10, ACM, pp. 1343–1352. doi:10.1145/1753326.1753529. 21
- [KMK18] KUCHER K., MARTINS R. M., KERREN A.: Analysis of VINCI 2009–2017 proceedings. In *Proceedings of the 11th International Symposium on Visual Information Communication and Interaction* (2018), VINCI '18, ACM, pp. 97–101. doi:10.1145/3231622.3231641. 19, 20
- [KMR17] KLEINBERG J., MULLAINATHAN S., RAGHAVAN M.: Inherent trade-offs in the fair determination of risk scores. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)* (2017), vol. 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 43:1–43:23. URL: <http://drops.dagstuhl.de/opus/volltexte/2017/8156>, doi:10.4230/LIPIcs.ITCS.2017.43. 7
- [Kos16] KOSARA R.: An empire built on sand: Reexamining what we think we know about visualization. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), BELIV '16, ACM, pp. 162–168. doi:10.1145/2993901.2993909. 28
- [KPB14] KRAUSE J., PERER A., BERTINI E.: INFUSE: Interactive feature selection for predictive modeling of high dimensional data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 1614–1623. doi:10.1109/TVCG.2014.2346482. 2, 15

- [KPB16] KRAUSE J., PERER A., BERTINI E.: Using visual analytics to interpret predictive machine learning models. In *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning* (2016), WHI '16. [arXiv:1606.05685](https://arxiv.org/abs/1606.05685). 17
- [KPB18] KRAUSE J., PERER A., BERTINI E.: A user study on the effect of aggregating explanations for interpreting machine learning models. In *Proceedings of the KDD Workshop on Interactive Data Exploration and Analytics* (2018), IDEA '18. URL: <http://poloclub.gatech.edu/idea2018/>. 13
- [KPHL16] KNUDSEN S., PEDERSEN J. G., HERDAL T., LARSEN J. E.: Using concrete and realistic data in evaluating initial visualization designs. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (New York, NY, USA, 2016), BELIV '16, ACM, pp. 27–35. [doi:10.1145/2993901.2993917](https://doi.org/10.1145/2993901.2993917). 28
- [KPK18] KUCHER K., PARADIS C., KERREN A.: The state of the art in sentiment visualization. *Computer Graphics Forum* 37, 1 (Feb. 2018), 71–96. [doi:10.1111/cgfm.13217](https://doi.org/10.1111/cgfm.13217). 13, 19, 27
- [KPN16] KRAUSE J., PERER A., NG K.: Interacting with predictions: Visual inspection of black-box machine learning models. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), CHI '16, ACM, pp. 5686–5697. [doi:10.1145/2858036.2858529](https://doi.org/10.1145/2858036.2858529). 19
- [KPSK17] KUCHER K., PARADIS C., SAHLGREN M., KERREN A.: Active learning and visual analytics for stance classification with ALVA. *ACM Transactions on Interactive Intelligent Systems* 7, 3 (Oct. 2017), 14:1–14:31. [doi:10.1145/3132169](https://doi.org/10.1145/3132169). 10
- [Kri09] KRIZHEVSKY A.: *Learning Multiple Layers of Features from Tiny Images*. Tech. rep., University of Toronto, 2009. 25, 26
- [KS12] KIENREICH W., SEIFERT C.: Visual exploration of feature-class matrices for classification problems. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2012), EuroVA '12, The Eurographics Association. [doi:10.2312/PE/EuroVAST/EuroVA12/037-041](https://doi.org/10.2312/PE/EuroVAST/EuroVA12/037-041). 17
- [KSH18] KARER B., SCHELER I., HAGEN H.: Panning for insight: Amplifying insight through tight integration of machine learning, data mining, and visualization. In *Proceedings of the EuroVis Workshop on Machine Learning Methods in Visualisation for Big Data* (2018), MLVis '18, The Eurographics Association. [doi:10.2312/mlvis.20181130](https://doi.org/10.2312/mlvis.20181130). 13
- [KTC*19] KAHNG M., THORAT N., CHAU D. H., VIÉGAS F. B., WATTENBERG M.: GAN Lab: Understanding complex deep generative models using interactive visual experimentation. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 310–320. [doi:10.1109/TVCG.2018.2864500](https://doi.org/10.1109/TVCG.2018.2864500). 18, 20
- [KZT*00] KEMP B., ZWINDERMAN A. H., TUK B., KAMPHUISEN H. A. C., OBERYE J. J. L.: Analysis of a sleep-dependent neuronal feedback loop: The slow-wave microcontinuity of the EEG. *IEEE Transactions on Biomedical Engineering* 47, 9 (Sept. 2000), 1185–1194. [doi:10.1109/10.867928](https://doi.org/10.1109/10.867928). 25
- [LA11] LESPINATS S., AUPETIT M.: CheckViz: Sanity check and topological clues for linear and non-linear mappings. *Computer Graphics Forum* 30, 1 (Mar. 2011), 113–125. [doi:10.1111/j.1467-8659.2010.01835.x](https://doi.org/10.1111/j.1467-8659.2010.01835.x). 7, 17, 27
- [Lan95] LANG K.: NewsWeeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on International Conference on Machine Learning* (1995), ICML '95, Morgan Kaufmann Publishers Inc., pp. 331–339. [doi:10.5555/3091622.3091662](https://doi.org/10.5555/3091622.3091662). 25, 26
- [LBBH98] LECUN Y., BOTTOU L., BENGIO Y., HAFFNER P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86, 11 (Nov. 1998), 2278–2324. [doi:10.1109/5.726791](https://doi.org/10.1109/5.726791). 25, 26
- [LBT*18] LIU S., BREMER P., THIAGARAJAN J. J., SRIKUMAR V., WANG B., LIVNAT Y., PASCUCCI V.: Visual exploration of semantic relationships in neural word embeddings. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 553–562. [doi:10.1109/TVCG.2017.2745141](https://doi.org/10.1109/TVCG.2017.2745141). 20
- [LCJ*18] LIU D., CUI W., JIN K., GUO Y., QU H.: DeepTracker: Visualizing the training process of convolutional neural networks. *ACM Transactions on Intelligent Systems and Technology* 10, 1 (Nov. 2018), 6:1–6:25. [doi:10.1145/3200489](https://doi.org/10.1145/3200489). 18
- [LCM*17] LU J., CHEN W., MA Y., KE J., LI Z., ZHANG F., MACIEJEWSKI R.: Recent progress and trends in predictive visual analytics. *Frontiers of Computer Science* 11, 2 (Apr. 2017), 192–207. [doi:10.1007/s11704-016-6028-y](https://doi.org/10.1007/s11704-016-6028-y). 8, 13
- [LGG*18] LIN H., GAO S., GOTZ D., DU F., HE J., CAO N.: RCLens: Interactive rare category exploration and identification. *IEEE Transactions on Visualization and Computer Graphics* 24, 7 (July 2018), 2223–2237. [doi:10.1109/TVCG.2017.2711030](https://doi.org/10.1109/TVCG.2017.2711030). 17, 22
- [LGH*17] LU Y., GARCIA R., HANSEN B., GLEICHER M., MACIEJEWSKI R.: The state-of-the-art in predictive visual analytics. *Computer Graphics Forum* 36, 3 (June 2017), 539–562. [doi:10.1111/cgfm.13210](https://doi.org/10.1111/cgfm.13210). 8, 9, 13, 17
- [LHF*18] LYONS J., HO N., FRIEDMAN J., ALARCON G., GUZNOV S.: Trust of learning systems: Considerations for code, algorithms, and affordances for learning. In *Human and Machine Learning: Visible, Explainable, Trustworthy and Transparent*. Springer International Publishing, 2018, pp. 265–278. [doi:10.1007/978-3-319-90403-0_13](https://doi.org/10.1007/978-3-319-90403-0_13). 3
- [Lin14] LIND N.: Better Life Index. In *Encyclopedia of Quality of Life and Well-Being Research*. Springer Netherlands, Dordrecht, 2014, pp. 381–382. [doi:10.1007/978-94-007-0753-5_3623](https://doi.org/10.1007/978-94-007-0753-5_3623). 25
- [LJLH19] LIU Y., JUN E., LI Q., HEER J.: Latent space cartography: Visual analysis of vector space embeddings. *Computer Graphics Forum* 38, 3 (June 2019), 67–78. [doi:10.1111/cgfm.13672](https://doi.org/10.1111/cgfm.13672). 20
- [LKC*12] LEE H., KIHM J., CHOO J., STASKO J., PARK H.: iVis-Clustering: An interactive visual document clustering via topic modeling. *Computer Graphics Forum* 31, 3pt3 (June 2012), 1155–1164. [doi:10.1111/j.1467-8659.2012.03108.x](https://doi.org/10.1111/j.1467-8659.2012.03108.x). 18
- [LKZ*15] LEHMANN D. J., KEMMLER F., ZHYHALAVA T., KIRSCHKE M., THEISEL H.: Visualnostics: Visual guidance pictograms for analyzing projections of high-dimensional data. *Computer Graphics Forum* 34, 3 (June 2015), 291–300. [doi:10.1111/cgfm.12641](https://doi.org/10.1111/cgfm.12641). 10
- [LL07] LENDASSE A., LIITAINEN E.: Variable scaling for time series prediction: Application to the ESTSP'07 and the NN3 forecasting competitions. In *Proceedings of the International Joint Conference on Neural Networks* (2007), IJCNN '07, IEEE, pp. 2812–2816. [doi:10.1109/IJCNN.2007.4371405](https://doi.org/10.1109/IJCNN.2007.4371405). 25
- [LL10] LASKOV P., LIPPMANN R.: Machine learning in adversarial environments. *Machine Learning* 81, 2 (Nov. 2010), 115–119. [doi:10.1007/s10994-010-5207-6](https://doi.org/10.1007/s10994-010-5207-6). 6
- [LLL*19] LIU S., LI Z., LI T., SRIKUMAR V., PASCUCCI V., BREMER P.: NLIZE: A perturbation-driven visual interrogation tool for analyzing and interpreting natural language inference models. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 651–660. [doi:10.1109/TVCG.2018.2865230](https://doi.org/10.1109/TVCG.2018.2865230). 17
- [LLS*18] LIU M., LIU S., SU H., CAO K., ZHU J.: Analyzing the noise robustness of deep neural networks. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2018), VAST '18, IEEE, pp. 60–71. [doi:10.1109/VAST.2018.8802509](https://doi.org/10.1109/VAST.2018.8802509). 18
- [LLWT15] LIU Z., LUO P., WANG X., TANG X.: Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), ICCV '15, IEEE, pp. 3730–3738. [doi:10.1109/ICCV.2015.425](https://doi.org/10.1109/ICCV.2015.425). 25
- [LMZ*14] LEE J. H., MCDONNELL K. T., ZELENYUK A., IMRE D., MUELLER K.: A structure-based distance metric for high-dimensional space exploration with multidimensional scaling. *IEEE Transactions on Visualization and Computer Graphics* 20, 3 (Mar. 2014), 351–364. [doi:10.1109/TVCG.2013.101](https://doi.org/10.1109/TVCG.2013.101). 21
- [LR02] LI X., ROTH D.: Learning question classifiers. In *Proceedings of the 19th International Conference on Computational Linguistics*

- tics — Volume 1 (2002), COLING '02, ACL, pp. 1–7. doi:10.3115/1072228.1072378. 25
- [LRL*18] LAUGEL T., RENARD X., LESOT M.-J., MARSALA C., DETYNECKI M.: Defining locality for surrogates in post-hoc interpretability. In *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning* (2018), WHI '18. arXiv: 1806.07498. 15
- [LS04] LEE J. D., SEE K. A.: Trust in automation: Designing for appropriate reliance. *Human Factors* 46, 1 (Mar. 2004), 50–80. doi: 10.1518/hfes.46.1.50_30392. 3
- [LSC*18] LIU M., SHI J., CAO K., ZHU J., LIU S.: Analyzing the training processes of deep generative models. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 77–87. doi: 10.1109/TVCG.2017.2744938. 17
- [LSL*17] LIU M., SHI J., LI Z., LI C., ZHU J., LIU S.: Towards better analysis of deep convolutional neural networks. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 91–100. doi: 10.1109/TVCG.2016.2598831. 16
- [LTBS*18] LÜCKE-TIEKE H., BEUTH M., SCHADER P., MAY T., BERNARD J., KOHLHAMMER J.: Lowering the barrier for successful replication and evaluation. In *Proceedings of the IEEE Workshop on Evaluation and Beyond — Methodological Approaches for Visualization* (2018), BELIV '18, IEEE, pp. 60–68. doi:10.1109/BELIV.2018.8634201. 25, 28
- [LWBP14] LIU S., WANG B., BREMER P.-T., PASCUCCI V.: Distortion-guided structure-driven interactive exploration of high-dimensional data. *Computer Graphics Forum* 33, 3 (June 2014), 101–110. doi: 10.1111/cgf.12366. 15
- [LWLZ17] LIU S., WANG X., LIU M., ZHU J.: Towards better analysis of machine learning models: A visual analytics perspective. *Visual Informatics* 1, 1 (Mar. 2017), 48–56. doi:10.1016/j.visinf.2017.01.006. 8, 13
- [LWT*15] LIU S., WANG B., THIAGARAJAN J. J., BREMER P.-T., PASCUCCI V.: Visual exploration of high-dimensional data through subspace analysis and dynamic projections. *Computer Graphics Forum* 34, 3 (June 2015), 271–280. doi:10.1111/cgf.12639. 18
- [LXL*18] LIU S., XIAO J., LIU J., WANG X., WU J., ZHU J.: Visual diagnosis of tree boosting methods. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 163–173. doi:10.1109/TVCG.2017.2744378. 17
- [Mad19] MADSEN A.: Visualizing memorization in RNNs. *Distill* (2019). doi:10.23915/distill.00016. 19
- [MAW19] MAWI working group traffic archive, 2019. Accessed January 10, 2020. URL: <https://mawi.wide.ad.jp/mawi/>. 25
- [MB02] MARTI U.-V., BUNKE H.: The IAM-database: An English sentence database for offline handwriting recognition. *International Journal on Document Analysis and Recognition* 5, 1 (Nov. 2002), 39–46. doi:10.1007/s100320200071. 25
- [MBD*11] MAY T., BANNACH A., DAVEY J., RUPPERT T., KOHLHAMMER J.: Guiding feature subset selection with an interactive visualization. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2011), VAST '11, IEEE, pp. 111–120. doi:10.1109/VAST.2011.6102448. 15
- [MBW11] MUNZNER T., BARSKY A., WILLIAMS M.: Reflections on QuestVis: A visualization system for an environmental sustainability model. In *Scientific Visualization: Interactions, Features, Metaphors* (2011), vol. 2 of *Dagstuhl Follow-Ups*, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 240–259. URL: <http://drops.dagstuhl.de/opus/volltexte/2011/3297>, doi: 10.4230/DFU.Vol2.SciViz.2011.240. 25
- [MCM*17] MA Y., CHEN W., MA X., XU J., HUANG X., MACIEJEWSKI R., TUNG A. K. H.: EasySVM: A visual analysis approach for open-box support vector machines. *Computational Visual Media* 3, 2 (2017), 161–175. doi:10.1007/s41095-017-0077-5. 16
- [MCMT14] MARTINS R. M., COIMBRA D. B., MINGHIM R., TELEA A. C.: Visual analysis of dimensionality reduction quality for parameterized projections. *Computers & Graphics* 41 (June 2014), 26–42. doi:10.1016/j.cag.2014.01.006. 16
- [MCR14] MORO S., CORTEZ P., RITA P.: A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems* 62 (June 2014), 22–31. doi:10.1016/j.dss.2014.03.001. 25
- [MCZ*17] MING Y., CAO S., ZHANG R., LI Z., CHEN Y., SONG Y., QU H.: Understanding hidden memories of recurrent neural networks. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2017), VAST '17, IEEE, pp. 13–24. doi: 10.1109/VAST.2017.8585721. 20
- [MDS95] MAYER R. C., DAVIS J. H., SCHOORMAN F. D.: An integrative model of organizational trust. *Academy of Management Review* 20, 3 (July 1995), 709–734. doi:10.5465/amr.1995.9508080335. 3
- [MHSW19] MAYR E., HYNEK N., SALISU S., WINDHAGER F.: Trust in information visualization. In *Proceedings of the EuroVis Workshop on Trustworthy Visualization* (2019), TrustVis '19, The Eurographics Association. doi:10.2312/trvis.201911187. 7
- [ML14] MOLCHANOV V., LINSEN L.: Interactive design of multidimensional data projection layout. In *Proceedings of the EG/VisGTC Conference on Visualization — Short Papers* (2014), EuroVis '14, The Eurographics Association. doi:10.2312/eurovisshort.20141152. 17
- [ML17] MCNABB L., LARAMEE R. S.: Survey of Surveys (SoS) — Mapping the landscape of survey papers in information visualization. *Computer Graphics Forum* 36, 3 (June 2017), 589–617. doi: 10.1111/cgf.13212. 8
- [MLMP18] MÜHLBACHER T., LINHARDT L., MÖLLER T., PIRINGER H.: TreePOD: Sensitivity-aware selection of pareto-optimal decision trees. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 174–183. doi:10.1109/TVCG.2017.2745158. 20
- [MMD*19] MURUGESAN S., MALIK S., DU F., KOH E., LAI T. M.: DeepCompare: Visual and interactive comparison of deep learning model performance. *IEEE Computer Graphics and Applications* 39, 5 (Sept. 2019), 47–59. doi:10.1109/MCG.2019.2919033. 17, 28
- [MMS93] MARCUS M. P., MARCINKIEWICZ M. A., SANTORINI B.: Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics* 19, 2 (June 1993), 313–330. doi:10.5555/972470.972475. 25
- [MP13] MÜHLBACHER T., PIRINGER H.: A partition-based framework for building and validating regression models. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec. 2013), 1962–1971. doi:10.1109/TVCG.2013.125. 19
- [MPG*14] MÜHLBACHER T., PIRINGER H., GRATZL S., SEDLMAIR M., STREIT M.: Opening the black box: Strategies for increased user involvement in existing algorithm implementations. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 1643–1652. doi:10.1109/TVCG.2014.2346578. 2
- [MQB19] MING Y., QU H., BERTINI E.: RuleMatrix: Visualizing and understanding classifiers with rules. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 342–352. doi:10.1109/TVCG.2018.2864812. 17
- [MRB*13] MANSOURI K., RINGSTED T., BALLABIO D., TODESCHINI R., CONSONNI V.: Quantitative structure–activity relationship models for ready biodegradability of chemicals. *Journal of Chemical Information and Modeling* 53, 4 (Apr. 2013), 867–878. doi:10.1021/ci4000213. 25
- [MRO*12] MACEACHREN A. M., ROTH R. E., O'BRIEN J., LI B., SWINGLEY D., GAHEGAN M.: Visual semiotics & uncertainty visualization: An empirical study. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2496–2505. doi:10.1109/TVCG.2012.279. 4
- [MSF*09] MATTHÄUS F., SMITH V. A., FOGTMAN A., SOMMER

- W. H., LEONARDI-ESSMANN F., LOURDUSAMY A., REIMERS M. A., SPANAGEL R., GEBICKE-HAERTER P. J.: Interactive molecular networks obtained by computer-aided conversion of microarray data from brains of alcohol-drinking rats. *Pharmacopsychiatry* 42 (May 2009), S118–S128. doi:10.1055/s-0029-1216348. 25
- [MSM*10] MEIRELLES P., SANTOS JR. C., MIRANDA J., KON F., TERCEIRO A., CHAVEZ C.: A study of the relationships between source code metrics and attractiveness in free software projects. In *Proceedings of the Brazilian Symposium on Software Engineering* (2010), SBES '10, IEEE, pp. 11–20. doi:10.1109/SBES.2010.27. 25
- [MSM*17] MICALLEF L., SUNDIN I., MARTTINEN P., AMMAD-UD DIN M., PELTOLA T., SOARE M., JACUCCI G., KASKI S.: Interactive elicitation of knowledge on feature relevance improves predictions in small data sets. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces* (2017), IUI '17, ACM, pp. 547–552. doi:10.1145/3025171.3025181. 13
- [MSSW16] MAYR E., SCHREDER G., SMUC M., WINDHAGER F.: Looking at the representations in our mind: Measuring mental models of information visualizations. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), BELIV '16, ACM, pp. 96–103. doi:10.1145/2993901.2993914. 28
- [MSW10] MACINNES J., SANTOSA S., WRIGHT W.: Visual classification: Expert knowledge guides machine learning. *IEEE Computer Graphics and Applications* 30, 1 (Jan. 2010), 8–14. doi:10.1109/MCG.2010.18. 16
- [MTCA17] MAGGIORI E., TARABALKA Y., CHARPIAT G., ALLIEZ P.: Can semantic labeling methods generalize to any city? The Inria Aerial Image Labeling Benchmark. In *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium* (2017), IGARSS '17, IEEE, pp. 3226–3229. doi:10.1109/IGARSS.2017.8127684. 25
- [Mun09] MUNZNER T.: A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (Nov. 2009), 921–928. doi:10.1109/TVCG.2009.111. 4, 6, 7
- [MvW11] MIGUT M. A., VAN GEMERT J. C., WORRING M.: Interactive decision making using dissimilarity to visually represented prototypes. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2011), VAST '11, IEEE, pp. 141–149. doi:10.1109/VAST.2011.6102451. 17
- [MW10] MIGUT M., WORRING M.: Visual exploration of classification models for risk assessment. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2010), VAST '10, IEEE, pp. 11–18. doi:10.1109/VAST.2010.5652398. 13
- [MXC*20] MING Y., XU P., CHENG F., QU H., REN L.: ProtoSteer: Steering deep sequence model with prototypes. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 238–248. doi:10.1109/TVCG.2019.2934267. 16
- [MXLM20] MA Y., XIE T., LI J., MACIEJEWSKI R.: Explaining vulnerabilities to adversarial machine learning through visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 1075–1085. doi:10.1109/TVCG.2019.2934631. 19, 28
- [MXQR19] MING Y., XU P., QU H., REN L.: Interpretable and steerable sequence learning via prototypes. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2019), KDD '19, ACM, pp. 903–913. doi:10.1145/3292500.3330908. 19, 28
- [MYA*13] MANDELLI D., YILMAZ A., ALDEMIR T., METZROTH K., DENNING R.: Scenario clustering and dynamic probabilistic risk assessment. *Reliability Engineering & System Safety* 115 (July 2013), 146–160. doi:10.1016/j.res.2013.02.013. 25
- [MYZ13] MIKOLOV T., YIH W.-T., ZWEIG G.: Linguistic regularities in continuous space word representations. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics* (2013), NAACL-HLT '13, ACL, pp. 746–751. URL: <https://aclweb.org/anthology/N13-1090>. 25
- [NA19] NONATO L. G., AUPETIT M.: Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE Transactions on Visualization and Computer Graphics* 25, 8 (Aug. 2019), 2650–2673. doi:10.1109/TVCG.2018.2846735. 9, 13
- [New17] New York Times articles, 2017. Accessed January 10, 2020. URL: <http://kaggle.com/nzalake52/new-york-times-articles>. 25
- [NGB*19] NALCACI A. A., GIRGIN D., BALKI S., TALAY F., BOZ H. A., BALCISOY S.: Detection of confirmation and distinction biases in visual analytics systems. In *Proceedings of the EuroVis Workshop on Trustworthy Visualization* (2019), TrustVis '19, The Eurographics Association. doi:10.2312/trvis.20191185. 7
- [NGDM*19] NIETO Y., GACÍA-DÍAZ V., MONTENEGRO C., GONZÁLEZ C. C., GONZÁLEZ CRESPO R.: Usage of machine learning for strategic decision making at higher educational institutions. *IEEE Access* 7 (2019), 75007–75017. doi:10.1109/ACCESS.2019.2919343. 1
- [NHM*07] NAM E. J., HAN Y., MUELLER K., ZELENYUK A., IMRE D.: ClusterSculptor: A visual analytics tool for high-dimensional data. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2007), VAST '07, IEEE, pp. 75–82. doi:10.1109/VAST.2007.4388999. 10
- [NHP*18] NIE S., HEALEY C., PADIA K., LEEMAN-MUNK S., BENSON J., CAIRA D., SETHI S., DEVARAJAN R.: Visualizing deep neural networks for text analytics. In *Proceedings of the IEEE Pacific Visualization Symposium* (2018), PacificVis '18, IEEE, pp. 180–189. doi:10.1109/PacificVis.2018.00031. 17
- [N.I17] N. I. P. Systems — NIPS 2017: Adversarial attacks and defenses, 2017. Accessed January 10, 2020. URL: <https://nips.cc/Conferences/2017/CompetitionTrack>. 25
- [NM13] NAM J. E., MUELLER K.: TripAdvisor^{N-D}: A tourism-inspired high-dimensional space exploration framework with overview and detail. *IEEE Transactions on Visualization and Computer Graphics* 19, 2 (Feb. 2013), 291–305. doi:10.1109/TVCG.2012.65. 25
- [NNM96] NENE S. A., NAYAR S. K., MURASE H.: *Columbia University Image Library (COIL-20)*. Tech. Rep. CUCS-005-96, Columbia University, Feb. 1996. URL: <http://cs.columbia.edu/CAVE/software/softlib/coil-20.php>. 25
- [NUM15] NUMBEO — Quality of life, 2015. Accessed January 10, 2020. URL: <https://numbeo.com/quality-of-life/>. 25
- [NW08] NORMAN M., WHALEN D.: IEEE Visualization 2008 Contest data, 2008. Accessed January 10, 2020. URL: <http://sciviscontest.ieeevis.org/2008/>. 25
- [NZ06] NILSBACK M.-E., ZISSERMAN A.: A visual vocabulary for flower classification. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2006), vol. 2 of CVPR '06, IEEE, pp. 1447–1454. doi:10.1109/CVPR.2006.42. 25
- [OAB*17] OLIVEIRA W., AMBRÓSIO L. M., BRAGA R., STRÖELE V., DAVID J. M., CAMPOS F.: A framework for provenance analysis and visualization. *Procedia Computer Science* 108 (2017), 1592–1601. doi:10.1016/j.procs.2017.05.216. 7
- [OOA10] OLUSOLA A. A., OLADELE A. S., ABOSEDE D. O.: Analysis of KDD '99 Intrusion Detection Dataset for selection of relevance features. In *Proceedings of the World Congress on Engineering and Computer Science* (2010), WCECS '10, International Association of Engineers, pp. 162–168. URL: <http://iaeng.org/publication/WCECS2010/>. 25
- [Ope14] OpenML — arsenic-female-bladder data set, 2014. Accessed January 10, 2020. URL: <https://openml.org/d/949>. 25
- [Ope19] Open Directory Project — Webpages and categories, 2019. Accessed January 10, 2020. URL: <https://dmoz-odp.org/>. 25
- [OSJ*18] OLAH C., SATYANARAYAN A., JOHNSON I., CARTER S., SCHUBERT L., YE K., MORDVINTSEV A.: The building blocks of interpretability. *Distill* (2018). doi:10.23915/distill.00010. 18

- [Ott14] Otto Group Product Classification Challenge, 2014. Accessed January 10, 2020. URL: <https://kaggle.com/c/otto-group-product-classification-challenge>. 25
- [OY03] OVER P., YEN J.: An introduction to DUC-2003: Intrinsic evaluation of generic news text summarization systems. In *Proceedings of the HLT 2003 Workshop on Text Summarization* (2003), DUC '03, NIST. URL: <https://duc.nist.gov/pubs.html#2003>. 25
- [PBK10] PIRINGER H., BERGER W., KRASSER J.: HyperMoVal: Interactive visual validation of regression models for real-time simulation. *Computer Graphics Forum* 29, 3 (June 2010), 983–992. doi: 10.1111/j.1467-8659.2009.01684.x. 15, 16
- [PCJB15] POZZOLO A. D., CAELEN O., JOHNSON R. A., BONTEMPI G.: Calibrating probability with undersampling for unbalanced classification. In *Proceedings of the IEEE Symposium Series on Computational Intelligence* (2015), SSCI '15, IEEE, pp. 159–166. doi: 10.1109/SSCI.2015.33. 25
- [PHL*16] PEZZOTTI N., HÖLLT T., LELIEVELDT B. P. F., EISEMANN E., VILANOVA A.: Hierarchical stochastic neighbor embedding. *Computer Graphics Forum* 35, 3 (June 2016), 21–30. doi:10.1111/cgf.12878. 11
- [PHV*18] PEZZOTTI N., HÖLLT T., VAN GEMERT J., LELIEVELDT B. P. F., EISEMANN E., VILANOVA A.: DeepEyes: Progressive visual analytics for designing deep neural networks. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 98–108. doi:10.1109/TVCG.2017.2744358. 16
- [PL05] PANG B., LEE L.: Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics* (2005), ACL '05, ACL, pp. 115–124. doi:10.3115/1219840.1219855. 25
- [PL08] PANG B., LEE L.: Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* 2, 1–2 (Jan. 2008), 1–135. doi:10.1561/1500000011. 25
- [PLHL19] PARK C., LEE J., HAN H., LEE K.: ComDia+: An interactive visual analytics system for comparing, diagnosing, and improving multiclass classifiers. In *Proceedings of the IEEE Pacific Visualization Symposium* (2019), PacificVis '19, IEEE, pp. 313–317. doi: 10.1109/PacificVis.2019.00044. 17
- [PLvdM*17] PEZZOTTI N., LELIEVELDT B. P. F., VAN DER MAATEN L., HÖLLT T., EISEMANN E., VILANOVA A.: Approximated and user steerable tSNE for progressive visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 23, 7 (July 2017), 1739–1752. doi:10.1109/TVCG.2016.2570755. 11
- [PNML08] PAULOVICH F. V., NONATO L. G., MINGHIM R., LEVKOWITZ H.: Least Square Projection: A fast high-precision multidimensional projection technique and its application to document mapping. *IEEE Transactions on Visualization and Computer Graphics* 14, 3 (May 2008), 564–575. doi:10.1109/TVCG.2007.70443. 25
- [PPM14] Parkinson's disease — Parkinsons Progression Markers Initiative (PPMI), 2014. Accessed January 10, 2020. URL: <http://www.ppmi-info.org/>. 25
- [PSF17] PELTONEN J., STRAHL J., FLORÉEN P.: Negative relevance feedback for exploratory search with visual interactive intent modeling. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces* (2017), IUI '17, ACM, pp. 149–159. doi: 10.1145/3025171.3025222. 17
- [PSMD14] PADUA L., SCHULZE H., MATKOVIĆ K., DELRIEUX C.: Interactive exploration of parameter space in data mining: Comprehending the predictive quality of large decision tree collections. *Computers & Graphics* 41 (June 2014), 99–113. doi:10.1016/j.cag.2014.02.004. 17
- [PSPM12] PAIVA J. G. S., SCHWARTZ W. R., PEDRINI H., MINGHIM R.: Semi-supervised dimensionality reduction based on Partial Least Squares for visual analysis of high dimensional data. *Computer Graphics Forum* 31, 3pt4 (June 2012), 1345–1354. doi:10.1111/j.1467-8659.2012.03126.x. 17, 25
- [PWJ06] PHILLIPS-WREN G., JAIN L.: Artificial intelligence for decision making. In *Knowledge-Based Intelligent Information and Engineering Systems (KES '06)* (2006), vol. 4251 of LNCS, Springer Berlin Heidelberg, pp. 531–536. doi:10.1007/11893004_69. 1
- [QH16] QU Z., HULLMAN J.: Evaluating visualization sets: Trade-offs between local effectiveness and global consistency. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), BELIV '16, ACM, pp. 44–52. doi: 10.1145/2993901.2993910. 28
- [RAL*17] REN D., AMERSHI S., LEE B., SUH J., WILLIAMS J. D.: Squares: Supporting interactive performance analysis for multiclass classifiers. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 61–70. doi:10.1109/TVCG.2016.2598828. 19
- [RB02] REDMOND M., BAVEJA A.: A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research* 141, 3 (Sept. 2002), 660–678. doi:10.1016/S0377-2217(01)00264-8. 25
- [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 Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 31–40. doi:10.1109/TVCG.2015.2467551. 7
- [Rev96] REVOW M.: Ringnorm Dataset, 1996. Accessed January 10, 2020. URL: <http://www.cs.toronto.edu/~delve/data/ringnorm/desc.html>. 25
- [RFFT17] RAUBER P. E., FADEL S. G., FALCÃO A. X., TELEA A. C.: Visualizing the hidden activity of artificial neural networks. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 101–110. doi:10.1109/TVCG.2016.2598838. 18
- [RFT18] RAUBER P. E., FALCÃO A. X., TELEA A. C.: Projections as visual aids for classification system design. *Information Visualization* 17, 4 (Oct. 2018), 282–305. doi:10.1177/1473871617713337. 17
- [RG19] ROESCH I., GÜNTHER T.: Visualization of neural network predictions for weather forecasting. *Computer Graphics Forum* 38, 1 (Feb. 2019), 209–220. doi:10.1111/cgf.13453. 16
- [RL14] RIECK B., LEITTE H.: Enhancing comparative model analysis using persistent homology. In *Proceedings of the IEEE VIS Workshop on Visualization for Predictive Analytics* (2014), VPA '14. URL: <http://predictive-workshop.github.io/>. 15
- [RL15a] RIECK B., LEITTE H.: Comparing dimensionality reduction methods using data descriptor landscapes. In *Proceedings of the Symposium on Visualization in Data Science at IEEE VIS* (2015), VDS '15. URL: <http://visualdatascience.org/2015/>. 18
- [RL15b] RIECK B., LEITTE H.: Persistent homology for the evaluation of dimensionality reduction schemes. *Computer Graphics Forum* 34, 3 (June 2015), 431–440. doi:10.1111/cgf.12655. 19
- [RL16] RIECK B., LEITTE H.: Exploring and comparing clusterings of multivariate data sets using persistent homology. *Computer Graphics Forum* 35, 3 (June 2016), 81–90. doi:10.1111/cgf.12884. 13
- [ŘS10] ŘEHŮŘEK R., SOJKA P.: Software framework for topic modelling with large corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks* (2010), ELRA, pp. 45–50. URL: <http://lrec-conf.org/proceedings/lrec2010/workshops/W10.pdf>. 20
- [RSF*15] RAUBER P. E., SILVA R. R. O. D., FERINGA S., CELEBI M. E., FALCÃO A. X., TELEA A. C.: Interactive image feature selection aided by dimensionality reduction. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2015), EuroVA '15, The Eurographics Association. doi:10.2312/eurova.20151098. 22
- [RSG16a] RIBEIRO M. T., SINGH S., GUESTRIN C.: Model-agnostic interpretability of machine learning. In *Proceedings of the ICML Workshop on Human Interpretability in Machine Learning* (2016), WHI '16. arXiv:1606.05386. 2

- [RSG16b] RIBEIRO M. T., SINGH S., GUESTRIN C.: “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016), KDD '16, ACM, pp. 1135–1144. doi:10.1145/2939672.2939778. 2
- [RSW02] ROSE T., STEVENSON M., WHITEHEAD M.: The Reuters Corpus volume 1 — From yesterday’s news to tomorrow’s language resources. In *Proceedings of the Third International Conference on Language Resources and Evaluation* (2002), LREC '02, ELRA. URL: <http://lrec-conf.org/proceedings/lrec2002/pdf/80.pdf>. 25
- [RU18] RUDIN C., USTUN B.: Optimized scoring systems: Toward trust in machine learning for healthcare and criminal justice. *INFORMS Journal on Applied Analytics* 48, 5 (Sept. 2018), 449–466. doi:10.1287/inte.2018.0957. 1
- [SAB*17] SEIFERT C., AAMIR A., BALAGOPALAN A., JAIN D., SHARMA A., GROTTTEL S., GUMHOLD S.: Visualizations of deep neural networks in computer vision: A survey. In *Transparent Data Mining for Big and Small Data*, vol. 32 of *Studies in Big Data*. Springer International Publishing, 2017, pp. 123–144. doi:10.1007/978-3-319-54024-5_6. 8, 13
- [SBE*18] SEVASTIANOVA R., BECK F., ELL B., TURKAY C., HENKIN R., BUTT M., KEIM D. A., EL-ASSADY M.: Going beyond visualization: Verbalization as complementary medium to explain machine learning models. In *Proceedings of the IEEE VIS Workshop on Visualization for AI Explainability* (2018), VISxAI '18. URL: <https://visxai.io/>. 17, 25, 27
- [SBIM12] SEDLMAIR M., BREHMER M., INGRAM S., MUNZNER T.: *Dimensionality Reduction in the Wild: Gaps and Guidance*. Tech. rep., Department of Computer Science, University of British Columbia, 2012. URL: <http://www.cs.ubc.ca/cgi-bin/tr/2012/TR-2012-03.2>
- [SBP19] SAWATZKY L., BERGNER S., POPOWICH F.: Visualizing RNN states with predictive semantic encodings. In *Proceedings of IEEE VIS 2019 — Short Papers* (2019), VIS '19, IEEE, pp. 156–160. doi:10.1109/VISUAL.2019.8933744. 20
- [SBTK08] SCHRECK T., BERNARD J., TEKUŠOVÁ T., KOHLHAMMER J.: Visual cluster analysis of trajectory data with interactive Kohonen maps. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2008), VAST '08, IEEE, pp. 3–10. doi:10.1109/VAST.2008.4677350. 10
- [Sch11] SCHULZ H.-J.: TreeVis.net: A tree visualization reference. *IEEE Computer Graphics and Applications* 31, 6 (Nov. 2011), 11–15. URL: <http://treevis.net>, doi:10.1109/MCG.2011.103. 27
- [Sco17] Scotch Whisky Dataset, 2017. Accessed January 10, 2020. URL: <https://kaggle.com/koki25ando/scotch-whisky-dataset>. 25
- [SDMT16] STAHNKE J., DÖRK M., MÜLLER B., THOM A.: Probing projections: Interaction techniques for interpreting arrangements and errors of dimensionality reductions. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 629–638. doi:10.1109/TVCG.2015.2467717. 16, 20
- [SDO19] Solar Dynamics Observatory (SDO), 2019. Accessed January 10, 2020. URL: <https://sdo.gsfc.nasa.gov/>. 25
- [SED*88] SMITH J. W., EVERHART J. E., DICKSON W., KNOWLER W. C., JOHANNES R. S.: Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. *Proceedings of the Annual Symposium on Computer Application in Medical Care* (Nov. 1988), 261–265. URL: <https://ncbi.nlm.nih.gov/pmc/articles/PMC2245318/>. 25
- [SGB*19] STROBELT H., GEHRMANN S., BEHRISCH M., PERER A., PFISTER H., RUSH A. M.: Seq2seq-Vis: A visual debugging tool for sequence-to-sequence models. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 353–363. doi:10.1109/TVCG.2018.2865044. 22
- [SGPR18] STROBELT H., GEHRMANN S., PFISTER H., RUSH A. M.: LSTMVis: A tool for visual analysis of hidden state dynamics in recurrent neural networks. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 667–676. doi:10.1109/TVCG.2017.2744158. 20
- [SGSG19] SIDEY-GIBBONS J. A. M., SIDEY-GIBBONS C. J.: Machine learning in medicine: A practical introduction. *BMC Medical Research Methodology* 19, 1 (2019), 64. 1
- [Shn00] SHNEIDERMAN B.: Designing trust into online experiences. *Communications of the ACM* 43, 12 (Dec. 2000), 57–59. doi:10.1145/355112.355124. 3
- [Shn20] SHNEIDERMAN B.: Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction* 36, 6 (2020), 495–504. doi:10.1080/10447318.2020.1741118. 3, 27
- [Sim12] SIMONOFF J. S.: *Smoothing Methods in Statistics*. Springer Series in Statistics. Springer-Verlag New York, New York, NY, USA, 2012. doi:10.1007/978-1-4612-4026-6. 25
- [SJS*17] SCHNEIDER B., JÄCKLE D., STOFFEL F., DIEHL A., FUCHS J., KEIM D. A.: Visual integration of data and model space in ensemble learning. In *Proceedings of the Symposium on Visualization in Data Science at IEEE VIS* (2017), VDS '17, IEEE, pp. 15–22. doi:10.1109/VDS.2017.8573444. 22
- [SJS*18] SCHNEIDER B., JÄCKLE D., STOFFEL F., DIEHL A., FUCHS J., KEIM D. A.: Integrating data and model space in ensemble learning by visual analytics. *IEEE Transactions on Big Data* (2018). doi:10.1109/TBDATA.2018.2877350. 16
- [SKB*18] SACHA D., KRAUS M., BERNARD J., BEHRISCH M., SCHRECK T., ASANO Y., KEIM D. A.: SOMFlow: Guided exploratory cluster analysis with self-organizing maps and analytic provenance. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 120–130. doi:10.1109/TVCG.2017.2744805. 18
- [SKBG*18] SMITH A., KUMAR V., BOYD-GRABER J., SEPPI K., FINDLATER L.: Closing the loop: User-centered design and evaluation of a human-in-the-loop topic modeling system. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces* (2018), IUI '18, ACM, pp. 293–304. doi:10.1145/3172944.3172965. 11
- [SKK*19] SHAH P., KENDALL F., KHOZIN S., GOOSEN R., HU J., LARAMIE J., RINGEL M., SCHORK N.: Artificial intelligence and machine learning in clinical development: A translational perspective. *npj Digital Medicine* 2, 1 (2019), 69. doi:10.1038/s41746-019-0148-3. 1
- [SKKC19] SACHA D., KRAUS M., KEIM D. A., CHEN M.: VIS4ML: An ontology for visual analytics assisted machine learning. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 385–395. doi:10.1109/TVCG.2018.2864838. 7, 16
- [SLT17] SUN Y., LANK E., TERRY M.: Label-and-learn: Visualizing the likelihood of machine learning classifier’s success during data labeling. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces* (2017), IUI '17, ACM, pp. 523–534. doi:10.1145/3025171.3025208. 6, 21
- [SMSL17] SHAO L., MAHAJAN A., SCHRECK T., LEHMANN D. J.: Interactive regression lens for exploring scatter plots. *Computer Graphics Forum* 36, 3 (June 2017), 157–166. doi:10.1111/cgf.13176. 13
- [SNLH09] SIPS M., NEUBERT B., LEWIS J. P., HANRAHAN P.: Selecting good views of high-dimensional data using class consistency. *Computer Graphics Forum* 28, 3 (June 2009), 831–838. doi:10.1111/j.1467-8659.2009.01467.x. 15
- [SNMM18] SHERKAT E., NOURASHRAFEDDIN S., MILIOS E. E., MINGHIM R.: Interactive document clustering revisited: A visual analytics approach. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces* (2018), IUI '18, ACM, pp. 281–292. doi:10.1145/3172944.3172964. 14
- [Sny19] SNYDER H.: Literature review as a research methodology: An

- overview and guidelines. *Journal of Business Research* 104 (2019), 333–339. doi:10.1016/j.jbusres.2019.07.039. 9
- [SPBA19] SALDANHA E., PRAGGASTIS B., BILLOW T., ARENDT D. L.: ReLVis: Visual analytics for situational awareness during reinforcement learning experimentation. In *Proceedings of the EG/VC Conference on Visualization — Short Papers* (2019), EuroVis '19, The Eurographics Association. doi:10.2312/evs.20191168. 16, 20, 28
- [SPG14] STOLPER C. D., PERER A., GOTZ D.: Progressive visual analytics: User-driven visual exploration of in-progress analytics. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 1653–1662. doi:10.1109/TVCG.2014.2346574. 13
- [Spi18] SPIEGELHALTER D.: Making algorithms trustworthy: What can statistical science contribute to transparency, explanation and validation? Plenary invited talk at NeurIPS '18, Dec. 2018. 4
- [Spo17] Spooky Author Identification, 2017. Accessed January 10, 2020. URL: <https://kaggle.com/bistaumanga/usps-dataset>. 25
- [SRG*18] SEHGAL G., RAWAT M., GUPTA B., GUPTA G., SHARMA G., SHROFF G.: Visual predictive analytics using iFuseML. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2018), EuroVA '18, The Eurographics Association. doi:10.2312/eurova.20181106. 16
- [SRM*15] SILVA R. R. O. D., RAUBER P. E., MARTINS R. M., MINGHIM R., TELEA A. C.: Attribute-based visual explanation of multidimensional projections. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2015), EuroVA '15, The Eurographics Association. doi:10.2312/eurova.20151100. 19
- [SSK10] SEIFERT C., SABOL V., KIENREICH W.: Stress Maps: Analysing local phenomena in dimensionality reduction based visualisations. In *Proceedings of the International Symposium on Visual Analytics Science and Technology* (2010), EuroVAST '10, The Eurographics Association. doi:10.2312/PE/EuroVAST/EuroVAST10/013-018. 19
- [SSK*16] SACHA D., SENARATNE H., KWON B. C., ELLIS G., KEIM D. A.: The role of uncertainty, awareness, and trust in visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 240–249. doi:10.1109/TVCG.2015.2467591. 2, 8
- [SSS*18] SCHALL M., SACHA D., STEIN M., FRANZ M. O., KEIM D. A.: Visualization-assisted development of deep learning models in offline handwriting recognition. In *Proceedings of the Symposium on Visualization in Data Science at IEEE VIS* (2018), VDS '18. URL: <http://visualdatascience.org/2018/>. 20
- [SSSEA20] SPINNER T., SCHLEGEL U., SCHÄFER H., EL-ASSADY M.: explAiner: A visual analytics framework for interactive and explainable machine learning. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 1064–1074. doi:10.1109/TVCG.2019.2934629. 19, 28
- [SSZ*16] SACHA D., SEDLMAIR M., ZHANG L., LEE J. A., WEISKOPF D., NORTH S. C., KEIM D. A.: Human-centered machine learning through interactive visualization: Review and open challenges. In *Proceedings of the 24th European Symposium on Artificial Neural Networks* (2016), ESANN 2016, Ciaco - i6doc.com, pp. 641–646. URL: <https://www.elen.ucl.ac.be/Proceedings/esann/esannpdf/es2016-166.pdf>. 8
- [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 (2017), 164–175. doi:10.1016/j.neucom.2017.01.105. 8
- [SvLB10] SCHRECK T., VON LANDESBERGER T., BREMM S.: Techniques for precision-based visual analysis of projected data. *Information Visualization* 9, 3 (Sept. 2010), 181–193. doi:10.1057/ivs.2010.2. 15
- [SW17] STREZOSKI G., WORRING M.: Plug-and-play interactive deep network visualization. In *Proceedings of the Workshop on Visual Analytics for Deep Learning* (2017), VADL '17. URL: <https://vadl2017.github.io/>. 13
- [SWM18] SAMEK W., WIEGAND T., MÜLLER K.-R.: Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *ITU Journal: ICT Discoveries* 1, 1 (Mar. 2018), 39–48. URL: <https://www.itu.int/en/journal/001/Pages/05.aspx>. 9
- [SYS*06] SMITH V. A., YU J., SMULDERS T. V., HARTEMINK A. J., JARVIS E. D.: Computational inference of neural information flow networks. *PLOS Computational Biology* 2, 11 (Nov. 2006). doi:10.1371/journal.pcbi.0020161. 25
- [SZD*00] SAMET J. M., ZEGER S. L., DOMINICI F., CURRIERO F., COURSAK I., DOCKERY D. W., SCHWARTZ J., ZANOBETTI A.: The National Morbidity, Mortality, and Air Pollution Study. Part II: Morbidity and mortality from air pollution in the United States. *Health Effects Institute Research Report*, 94 part II (June 2000), 5–70. URL: <https://healtheffects.org/publication/national-morbidity-mortality-and-air-pollution-study-part-ii-morbidity-and-mortality-air>. 25
- [SZK16] SMARR B. L., ZUCKER I., KRIEGSFELD L. J.: Detection of successful and unsuccessful pregnancies in mice within hours of pairing through frequency analysis of high temporal resolution core body temperature data. *PLOS ONE* 11, 7 (July 2016). doi:10.1371/journal.pone.0160127. 25
- [SZL*18] SUN J., ZHU Q., LIU Z., LIU X., LEE J., SU Z., SHI L., HUANG L., XU W.: FraudVis: Understanding unsupervised fraud detection algorithms. In *Proceedings of the IEEE Pacific Visualization Symposium* (2018), PacificVis '18, IEEE, pp. 170–174. doi:10.1109/PacificVis.2018.00029. 20
- [SZS*17] SACHA D., ZHANG L., SEDLMAIR M., LEE J. A., PELTONEN J., WEISKOPF D., NORTH S. C., KEIM D. A.: Visual interaction with dimensionality reduction: A structured literature analysis. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 241–250. doi:10.1109/TVCG.2016.2598495. 9
- [TA13] TOMINSKI C., AIGNER W.: The TimeVis Browser, 2013. Accessed January 10, 2020. URL: <http://survey.timeviz.net>. 27
- [TAC*20] TOREINI E., AITKEN M., COOPAMOOTOO K., ELLIOTT K., ZELAYA C. G., VAN MOORSEL A.: The relationship between trust in AI and trustworthy machine learning technologies. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (2020), FAT* '20, ACM, pp. 272–283. doi:10.1145/3351095.3372834. 1
- [Tay90] TAYLOR R.: Interpretation of the correlation coefficient: A basic review. *Journal of Diagnostic Medical Sonography* 6, 1 (Jan. 1990), 35–39. doi:10.1177/875647939000600106. 22
- [TCE*19] TYAGI A., CAO Z., ESTRO T., ZADOK E., MUELLER K.: ICE: An interactive configuration explorer for high dimensional categorical parameter spaces. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2019), VAST '19, IEEE. arXiv:1907.12627. 20
- [Tit15] Titanic: Machine learning from disaster, 2015. Accessed January 10, 2020. URL: <https://kaggle.com/c/titanic>. 25
- [TKDB17] TAMAGNINI P., KRAUSE J., DASGUPTA A., BERTINI E.: Interpreting black-box classifiers using instance-level visual explanations. In *Proceedings of the 2nd Workshop on Human-In-the-Loop Data Analytics* (2017), HILDA '17, ACM, pp. 6:1–6:6. doi:10.1145/3077257.3077260. 21
- [TKK18] THAMMACHANTUEK I., KOSOLSOMNAT S., KETCHAM M.: Comparison of machine learning algorithm's performance based on decision making in autonomous car. In *Proceedings of the International Joint Symposium on Artificial Intelligence and Natural Language Processing* (2018), iSAI-NLP '18, IEEE. doi:10.1109/iSAI-NLP.2018.8693002. 1

- [TLKT09] TALBOT J., LEE B., KAPOOR A., TAN D. S.: EnsembleMatrix: Interactive visualization to support machine learning with multiple classifiers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2009), CHI '09, ACM, pp. 1283–1292. doi:10.1145/1518701.1518895. 16
- [TLRB18] THIAGARAJAN J. J., LIU S., RAMAMURTHY K. N., BREMER P.-T.: Exploring high-dimensional structure via axis-aligned decomposition of linear projections. *Computer Graphics Forum* 37, 3 (June 2018), 241–251. doi:10.1111/cgfm.13416. 19
- [TMF*12] TATU A., MAASS F., FÄRBER I., BERTINI E., SCHRECK T., SEIDL T., KEIM D. A.: Subspace search and visualization to make sense of alternative clusterings in high-dimensional data. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2012), VAST '12, IEEE, pp. 63–72. doi:10.1109/VAST.2012.6400488. 15
- [TPRH11] TURKAY C., PARULEK J., REUTER N., HAUSER H.: Interactive visual analysis of temporal cluster structures. *Computer Graphics Forum* 30, 3 (June 2011), 711–720. doi:10.1111/j.1467-8659.2011.01920.x. 16
- [Tra19] TransitFeeds — Nashville MTA GTFS, 2019. Accessed January 10, 2020. URL: <https://transitfeeds.com/p/nashville-mta/220>. 25
- [TSL00] TENENBAUM J. B., SILVA V. D., LANGFORD J. C.: A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 5500 (Dec. 2000), 2319–2323. doi:10.1126/science.290.5500.2319. 25
- [TSL*16] TURKAY C., SLINGSBY A., LAHTINEN K., BUTT S., DYKES J.: Enhancing a social science model-building workflow with interactive visualisation. In *Proceedings of the 24th European Symposium on Artificial Neural Networks* (2016), ESANN 2016, Ciaco - i6doc.com, pp. 629–634. URL: <http://www.elen.ucl.ac.be/Proceedings/esann/esannpdf/es2016-147.pdf>. 15
- [TTKV01] TETKO I. V., TANCHUK V. Y., KASHEVA T. N., VILLA A. E. P.: Estimation of aqueous solubility of chemical compounds using E-state indices. *Journal of Chemical Information and Computer Sciences* 41, 6 (Nov. 2001), 1488–1493. doi:10.1021/ci000392t. 25
- [TV07] TRON R., VIDAL R.: A benchmark for the comparison of 3-D motion segmentation algorithms. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2007), CVPR '07, IEEE. doi:10.1109/CVPR.2007.382974. 25
- [TZY*08] TANG J., ZHANG J., YAO L., LI J., ZHANG L., SU Z.: ArnetMiner: Extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2008), KDD '08, ACM, pp. 990–998. doi:10.1145/1401890.1402008. 25
- [USD19] USDA National Nutrient Database, 2019. Accessed January 10, 2020. URL: <https://fdc.nal.usda.gov/>. 25, 26
- [USP17] Handwritten Digits USPS Dataset, 2017. Accessed January 10, 2020. URL: <https://kaggle.com/bistaumanga/usps-dataset>. 25
- [vdMH08] VAN DER MAATEN L., HINTON G.: Visualizing data using t-SNE. *Journal of Machine Learning Research* 9 (2008), 2579–2605. URL: <http://jmlr.org/papers/v9/vandermaaten08a.html>. 22
- [VDMPvdH09] VAN DER MAATEN L., POSTMA E., VAN DEN HERIK J.: Dimensionality reduction: a comparative review. *Journal of Machine Learning Research* 10 (2009), 66–71. 8, 9
- [vdPvS00] VAN DER PUTTEN P., VAN SOMEREN M.: *CoIL Challenge 2000: The Insurance Company Case*. Tech. Rep. 2000-09, Leiden Institute of Advanced Computer Science, 2000. URL: <http://liacs.leidenuniv.nl/~puttenpwhvander/library/cc2000/>. 25
- [VKA*18] VOGOGIAS A., KENNEDY J., ARCHAMBAULT D., BACH B., SMITH V. A., CURRANT H.: BayesPiles: Visualisation support for bayesian network structure learning. *ACM Transactions on Intelligent Systems and Technology* 10, 1 (Nov. 2018), 5:1–5:23. doi:10.1145/3230623. 20
- [VL09] VAN LONG T., LINSEN L.: MultiClusterTree: Interactive visual exploration of hierarchical clusters in multidimensional multivariate data. *Computer Graphics Forum* 28, 3 (June 2009), 823–830. doi:10.1111/j.1467-8659.2009.01468.x. 16
- [VSK*15] VIECHTBAUER W., SMITS L., KOTZ D., BUDÉ L., SPIGT M., SERROYEN J., CRUTZEN R.: A simple formula for the calculation of sample size in pilot studies. *Journal of Clinical Epidemiology* 68, 11 (2015), 1375–1379. doi:10.1016/j.jclinepi.2015.04.014. 9
- [vULM*16] VAN UNEN V., LI N., MOLENDIJK I., TEMURHAN M., HÖLLT T., VAN DER MEULEN-DE JONG A., VERSPAGET H., MEARIN M., MULDER C., VAN BERGEN J., LELIEVELDT B., KONING F.: Mass cytometry of the human mucosal immune system identifies tissue- and disease-associated immune subsets. *Immunity* 44, 5 (May 2016), 1227–1239. doi:10.1016/j.immuni.2016.04.014. 25
- [vv11] VAN DEN ELZEN S., VAN WIJK J. J.: BaobabView: Interactive construction and analysis of decision trees. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2011), VAST '11, IEEE, pp. 151–160. doi:10.1109/VAST.2011.6102453. 17, 20
- [VvDW*15] VETA M., VAN DIEST P. J., WILLEMS S. M., WANG H., MADABHUSHI A., CRUZ-ROA A., GONZALEZ F., LARSEN A. B., VESTERGAARD J. S., DAHL A. B., CIREŞAN D. C., SCHMIDHUBER J., GIUSTI A., GAMBARDILLA L. M., TEK F. B., WALTER T., WANG C.-W., KONDO S., MATUSZEWSKI B. J., PRECIOSO F., SNELL V., KITTLER J., DE CAMPOS T. E., KHAN A. M., RAJPOOT N. M., ARKOUMANI E., LACLE M. M., VIERGEVER M. A., PLUIM J. P.: Assessment of algorithms for mitosis detection in breast cancer histopathology images. *Medical Image Analysis* 20, 1 (Feb. 2015), 237–248. doi:10.1016/j.media.2014.11.010. 25
- [WCH*15] WU F., CHEN G., HUANG J., TAO Y., CHEN W.: EasyXplorer: A flexible visual exploration approach for multivariate spatial data. *Computer Graphics Forum* 34, 7 (Oct. 2015), 163–172. doi:10.1111/cgfm.12755. 10
- [WCR*18] WENSKOVITCH J., CRANDELL I., RAMAKRISHNAN N., HOUSE L., LEMAN S., NORTH C.: Towards a systematic combination of dimension reduction and clustering in visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 131–141. doi:10.1109/TVCG.2017.2745258. 10
- [WGG10] WU Y., GAUNT C., GRAY S.: A comparison of alternative bankruptcy prediction models. *Journal of Contemporary Accounting & Economics* 6, 1 (June 2010), 34–45. doi:10.1016/j.jcae.2010.04.002. 25
- [WGSY19] WANG J., GOU L., SHEN H., YANG H.: DQNviz: A visual analytics approach to understand Deep Q-Networks. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 288–298. doi:10.1109/TVCG.2018.2864504. 17
- [WGYS18] WANG J., GOU L., YANG H., SHEN H.: GANviz: A visual analytics approach to understand the adversarial game. *IEEE Transactions on Visualization and Computer Graphics* 24, 6 (June 2018), 1905–1917. doi:10.1109/TVCG.2018.2816223. 17
- [WGZ*19] WANG J., GOU L., ZHANG W., YANG H., SHEN H.: DeepVID: Deep visual interpretation and diagnosis for image classifiers via knowledge distillation. *IEEE Transactions on Visualization and Computer Graphics* 25, 6 (June 2019), 2168–2180. doi:10.1109/TVCG.2019.2903943. 13, 14, 20
- [WHO19] World Health Organization (WHO-SIS) Statistical Information System, 2019. Accessed January 10, 2020. URL: <http://who.int/whosis/en/>. 25
- [WJCC16] WHITEHEAD A. L., JULIOUS S. A., COOPER C. L., CAMPBELL M. J.: Estimating the sample size for a pilot randomised trial to minimise the overall trial sample size for the external pilot and main trial for a continuous outcome variable. *Statistical Methods in*

- Medical Research* 25, 3 (June 2016), 1057–1073. doi:10.1177/0962280215588241. 9
- [WLN*17] WANG Y., LI J., NIE F., THEISEL H., GONG M., LEHMANN D. J.: Linear discriminative star coordinates for exploring class and cluster separation of high dimensional data. *Computer Graphics Forum* 36, 3 (June 2017), 401–410. doi:10.1111/cgf.13197. 19
- [WLS19] WANG J., LIU X., SHEN H.-W.: High-dimensional data analysis with subspace comparison using matrix visualization. *Information Visualization* 18, 1 (Jan. 2019), 94–109. doi:10.1177/1473871617733996. 22
- [WM18] WANG B., MUELLER K.: The Subspace Voyager: Exploring high-dimensional data along a continuum of salient 3D subspaces. *IEEE Transactions on Visualization and Computer Graphics* 24, 2 (Feb. 2018), 1204–1222. doi:10.1109/TVCG.2017.2672987. 15
- [WMJ*19] WANG Q., MING Y., JIN Z., SHEN Q., LIU D., SMITH M. J., VEERAMACHANENI K., QU H.: ATMSeer: Increasing transparency and controllability in automated machine learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), CHI '19, ACM, pp. 681:1–681:12. doi:10.1145/3290605.3300911. 16
- [Woh14] WOHLIN C.: Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering* (2014), EASE '14, ACM, pp. 38:1–38:10. doi:10.1145/2601248.2601268. 10
- [Wol19] WOLF C. T.: Explainability scenarios: Towards scenario-based XAI design. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (2019), IUI '19, ACM, pp. 252–257. doi:10.1145/3301275.3302317. 2
- [WPB*20] 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* 26, 1 (Jan. 2020), 56–65. doi:10.1109/TVCG.2019.2934619. 17, 18
- [WSW*18] WONGSUPHASAWAT K., SMILKOV D., WEXLER J., WILSON J., MANÉ D., FRITZ D., KRISHNAN D., VIÉGAS F. B., WATTENBERG M.: Visualizing dataflow graphs of deep learning models in TensorFlow. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 1–12. doi:10.1109/TVCG.2017.2744878. 19
- [WvP11] WIERING M. A., VAN HASSELT H., PIETERSMA A.-D., SCHOMAKER L.: Reinforcement learning algorithms for solving classification problems. In *Proceedings of the IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning* (2011), ADPRL '11, IEEE, pp. 91–96. doi:10.1109/ADPRL.2011.5967372. 28
- [XCH*16] XIA J., CHEN W., HOU Y., HU W., HUANG X., EBERTK D. S.: DimScanner: A relation-based visual exploration approach towards data dimension inspection. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (2016), VAST '16, IEEE, pp. 81–90. doi:10.1109/VAST.2016.7883514. 17
- [XMRC17] XU P., MEI H., REN L., CHEN W.: ViDX: Visual diagnostics of assembly line performance in smart factories. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan. 2017), 291–300. doi:10.1109/TVCG.2016.2598664. 25
- [XPGF19] XIONG C., PADILLA L., GRAYSON K., FRANCONERI S.: Examining the components of trust in map-based visualizations. In *Proceedings of the EuroVis Workshop on Trustworthy Visualization* (2019), TrustVis '19, The Eurographics Association. doi:10.2312/trvis.20191186. 7
- [XRV17] XIAO H., RASUL K., VOLLGRAF R.: Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms, 2017. arXiv:1708.07747. 25
- [XXM*19] XU K., XIA M., MU X., WANG Y., CAO N.: EnsembleLens: Ensemble-based visual exploration of anomaly detection algorithms with multidimensional data. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 109–119. doi:10.1109/TVCG.2018.2864825. 22
- [XYC*18] XIA J., YE F., CHEN W., WANG Y., CHEN W., MA Y., TUNG A. K. H.: LDSScanner: Exploratory analysis of low-dimensional structures in high-dimensional datasets. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 236–245. doi:10.1109/TVCG.2017.2744098. 17
- [YBC*18] YU K., BERKOVSKY S., CONWAY D., TAIB R., ZHOU J., CHEN F.: Do I trust a machine? Differences in user trust based on system performance. In *Human and Machine Learning: Visible, Explainable, Trustworthy and Transparent*, HCIS. Springer International Publishing, 2018, pp. 245–264. doi:10.1007/978-3-319-90403-0_12. 3
- [YCN*15] YOSINSKI J., CLUNE J., NGUYEN A., FUCHS T., LIPSON H.: Understanding neural networks through deep visualization. In *Proceedings of the ICML Workshop on Visualization for Deep Learning* (2015), DL '15. arXiv:1506.06579. 20
- [YDC*12] YAN S., DONG J., CHEN Q., SONG Z., PAN Y., XIA W., HUANG Z., HUA Y., SHEN S.: Generalized hierarchical matching for sub-category aware object classification. In *Proceedings of the ECCV 2012 PASCAL Visual Object Classes Challenge Workshop* (2012), VOC '12. URL: <http://host.robots.ox.ac.uk:8080/pascal/VOC/voc2012/workshop/index.html>. 25
- [Yeh98] YE H. I.-C.: Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research* 28, 12 (Dec. 1998), 1797–1808. doi:10.1016/S0008-8846(98)00165-3. 25
- [Yel19] Yelp Open Dataset, 2019. Accessed January 10, 2020. URL: <https://yelp.com/dataset/>. 25
- [YL09] YE H. I.-C., LIEN C.-H.: The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications* 36, 2 (Mar. 2009), 2473–2480. doi:10.1016/j.eswa.2007.12.020. 25
- [YRW07] YANG D., RUNDENSTEINER E. A., WARD M. O.: Analysis guided visual exploration of multivariate data. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology* (2007), VAST '07, IEEE, pp. 83–90. doi:10.1109/VAST.2007.4389000. 25
- [YS18] YU R., SHI L.: A user-based taxonomy for deep learning visualization. *Visual Informatics* 2, 3 (Sept. 2018), 147–154. doi:10.1016/j.visinf.2018.09.001. 9, 13
- [YZR*18] YAN L., ZHAO Y., ROSEN P., SCHEIDEGGER C., WANG B.: Homology-preserving dimensionality reduction via manifold landmarking and tearing. In *Proceedings of the Symposium on Visualization in Data Science at IEEE VIS* (2018), VDS '18. URL: <http://visualdatascience.org/2018/>. 19
- [ZCW*19] ZHAO X., CUI W., WU Y., ZHANG H., QU H., ZHANG D.: Oui! Outlier interpretation on multi-dimensional data via visual analytics. *Computer Graphics Forum* 38, 3 (June 2019), 213–224. doi:10.1111/cgf.13683. 18
- [ZF14] ZEILER M. D., FERGUS R.: Visualizing and understanding convolutional networks. In *Proceedings of the 13th European Conference on Computer Vision (ECCV '14)* (2014), vol. 8689 of LNCS, Springer International Publishing, pp. 818–833. doi:10.1007/978-3-319-10590-1_53. 17
- [ZHP*17] ZENG H., HALEEM H., PLANTAZ X., CAO N., QU H.: CN-Comparator: Comparative analytics of convolutional neural networks. In *Proceedings of the Workshop on Visual Analytics for Deep Learning* (2017), VADL '17. URL: <https://vadl2017.github.io/>. 15
- [ZI05] ZELENYUK A., IMRE D.: Single Particle Laser Ablation Time-of-Flight Mass Spectrometer: An introduction to SPLAT. *Aerosol Science and Technology* 39, 6 (June 2005), 554–568. doi:10.1080/027868291009242. 25

- [ZKM*19] ZHAO J., KARIMZADEH M., MASJEDI A., WANG T., ZHANG X., CRAWFORD M. M., EBERT D. S.: FeatureExplorer: Interactive feature selection and exploration of regression models for hyperspectral images. In *Proceedings of IEEE VIS 2019 — Short Papers* (2019), VIS '19, IEEE, pp. 161–165. doi:10.1109/VISUAL.2019.8933619. 15
- [ZLH*16] ZHOU F., LI J., HUANG W., ZHAO Y., YUAN X., LIANG X., SHI Y.: Dimension reconstruction for visual exploration of subspace clusters in high-dimensional data. In *Proceedings of the IEEE Pacific Visualization Symposium* (2016), PacificVis '16, IEEE, pp. 128–135. doi:10.1109/PACIFICVIS.2016.7465260. 15
- [ZSCC18] ZHAO J., SUN M., CHEN F., CHIU P.: BiDots: Visual exploration of weighted biclusters. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 195–204. doi:10.1109/TVCG.2017.2744458. 10
- [ZTR16] ZHAO Y., TASOULIS S., ROOS T.: Manifold visualization via short walks. In *Proceedings of the EG/VGTC Conference on Visualization — Short Papers* (2016), EuroVis '16, Eurographics Association, pp. 85–89. doi:10.2312/eurovisshort.20161166. 17
- [ZWC*18] ZHAO X., WU Y., CUI W., DU X., CHEN Y., WANG Y., LEE D. L., QU H.: SkyLens: Visual analysis of skyline on multi-dimensional data. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 246–255. doi:10.1109/TVCG.2017.2744738. 17
- [ZWLC19] ZHAO X., WU Y., LEE D. L., CUI W.: iForest: Interpreting random forests via visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 407–416. doi:10.1109/TVCG.2018.2864475. 15
- [ZWM*19] ZHANG J., WANG Y., MOLINO P., LI L., EBERT D. S.: Manifold: A model-agnostic framework for interpretation and diagnosis of machine learning models. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 364–373. doi:10.1109/TVCG.2018.2864499. 18
- [ZWRH14] ZHAO K., WARD M. O., RUNDENSTEINER E. A., HIGGINS H. N.: LoVis: Local pattern visualization for model refinement. *Computer Graphics Forum* 33, 3 (June 2014), 331–340. doi:10.1111/cgf.12389. 18, 28
- [ZYB*16] ZHANG C., YANG J., BENJAMIN ZHAN F., GONG X., BRENDER J. D., LANGLOIS P. H., BARLOWE S., ZHAO Y.: A visual analytics approach to high-dimensional logistic regression modeling and its application to an environmental health study. In *Proceedings of the IEEE Pacific Visualization Symposium* (2016), PacificVis '16, IEEE, pp. 136–143. doi:10.1109/PACIFICVIS.2016.7465261. 21
- [ZZ18] ZHANG Q.-S., ZHU S.-C.: Visual interpretability for deep learning: A survey. *Frontiers of Information Technology & Electronic Engineering* 19, 1 (Jan. 2018), 27–39. doi:10.1631/FITEE.1700808. 9