Computational Visual Media
DOI 10.1007/s41095-xxx-xxxx-x Vol. x, No. x, month year, xx–xx

Article Type (Research/Review)

A Survey of Visual Analytics Techniques for Machine Learning

Jun Yuan[1], Changjian Chen[1], Weikai Yang[1], Mengchen Liu[2], Jiazhi Xia[3], Shixia Liu[1](✉)

© The Author(s) 2020. This article is published with open access at Springerlink.com

Abstract Visual analytics for machine learning has recently evolved as one of the most exciting areas in the field of visualization. To better identify which research topics are promising and learn how to apply relevant techniques in visual analytics applications, we systematically review 259 papers that are published in the recent ten years or the representative works before 2010. We build a taxonomy, which includes three first-level categories: techniques before model building, techniques in modeling building, and techniques after model building. Each category is further characterized by representative analysis tasks, and each task is exemplified by a set of recent influential works. We also discuss and highlight research challenges and potential future research opportunities that can be promising and useful for visual analytics researchers.

Keywords visual analytics; machine learning; data quality; feature selection; model understanding; content analysis.

1 Introduction

The recent success of artificial intelligence applications depends on the performance and capabilities of machine learning models [162]. In the past ten years, a variety of visual analytics methods have been proposed to make machine learning more explainable, trustworthy, and reliable. These research efforts fully combine the advantages of interactive visualization and machine learning techniques to facilitate the analysis and understanding of the major components in the learning process, with an aim to improve performance. For example, visual analytics research for explaining the inner workings of deep convolutional neural networks has increased the transparency of deep learning models and received continuous, even more enormous attention recently [54, 103, 162, 285].

The rapid development of visual analytics techniques for machine learning yields an emerging need for a comprehensive analysis of this area to support the understanding of how visualization techniques are designed and applied in machine learning pipelines. There are several initial efforts to summarize the advances in this field from different aspects. For example, Liu *et al.* [161] summarized the visualization techniques for text analysis. Lu *et al.* [172] surveyed visual analytics techniques for predictive models. Recently, Liu *et al.* [162] presented a viewpoint paper on the analysis of machine learning models from the visual analytics perspective. Sacha *et al.* [217] analyzed a set of example systems and proposed an ontology for visual analytics assisted machine learning. However, existing surveys either focus on a specific area of machine learning (*e.g.*, text mining [161], predictive model [172], model understanding [162]) or aim to sketch an ontology [217] based on a set of example techniques only.

In this paper, we aim to provide a comprehensive survey of visual analytics techniques for machine learning, which focuses on every phase of the machine learning pipeline. We focus on the works in the visualization community. Still, the AI community has also made solid contributions to the study of visually explaining the feature detectors in deep learning

1 BNRist, Tsinghua University, Beijing 100086, China. Email: J. Yuan, C. Chen, W. Yang, {yuanj19, ccj17, ywk19}@mails.tsinghua.edu.cn; S. Liu, shixia@tsinghua.edu.cn(✉).

2 Microsoft, Redmond 98052, United States, E-mail: mengcliu@microsoft.com.

3 Central South University, Changsha 410083, China, E-mail: xiajiazhi@csu.edu.cn.

Manuscript received: xxxx-xx-xx; accepted: xxxx-xx-xx.

models. For example, Selvaraju *et al.* [221] tried to identify the sensitive part of an image to its classification result by computing the class activation maps. Readers can refer to the surveys of Zhang *et al.* [226] and Hohman *et al.* [103] for more details. We collect 259 papers from related top-tier venues in the past ten years through a systematical procedure. According to the machine learning pipeline, we divide the literature into three stages: before, in, and after model building. We analyze the functions of visual analytics techniques in the three stages and abstract typical tasks, including improving data quality and feature quality before model building, model understanding, diagnosing, and steering in modeling building, data understanding after model building. Each task is featured by a set of carefully selected examples. We highlight six prominent research directions and open problems in the field of visual analytics for machine learning. We hope that this survey acts as the context to discuss machine learning related visual analytics techniques and a start point for practitioners and researchers to develop visual analytics tools for machine learning.

2 Survey Landscape

2.1 Paper Selection

In this paper, we focus on visual analytics techniques that help develop explainable, trustworthy, and reliable machine learning applications. To comprehensively survey visual analytics techniques for machine learning, we performed an exhausting manual review of related top-tier venues in the past ten years (2010-2020). The venues consist InfoVis, VAST, Vis (later SciVis), EuroVis, PacificVis, IEEE TVCG, CGF, and CG&A. The manual review was conducted by three Ph.D. candidates with more than two years of research experience in visual analytics. We followed the manual review process in the text visualization survey [161]. Specifically, we first reviewed the titles of papers from the aforementioned venues to identify candidate papers. Next, we reviewed the abstracts of identified candidate papers to further determine whether they are visual analytics techniques for machine learning. If the title and abstract cannot provide clear information, the full text was gone through for the final decision. In addition to the exhausting manual review of the above venues, we also searched for the representative related works that appeared previously or in other venues, such as the Profiler [123].

After the reviewing and searching process, 259 papers are selected. Tab. 1 presents detailed statistics. Due to the boosting of machine learning techniques in the past ten years, this field has been attracting more and more research attention.

2.2 Taxonomy

In this section, we comprehensively analyze the collected visual analytics works for systematically understanding the major research trends. These works are categorized based on a typical machine learning pipeline [182] to solve real-world problems. As shown in Fig. 1, such a pipeline contains three stages: (1) data pre-processing before model building, (2) machine learning model building, and (3) deployment after the model is built. Accordingly, the visual analytics techniques for machine learning can be mapped into these three stages: techniques before model building, techniques in model building, and techniques after model building.

2.2.1 Techniques before Model Building

The major goal of visual analytics techniques before model building is to help model developers better prepare the data for model building. The quality of the data is mainly determined by the data itself and the features we use. Accordingly, there are two research directions, i.e., visual analytics for data quality improvement and feature engineering.

Data quality can be improved from various aspects, such as completing the missing data attributes and correcting wrong data labels. Previously, these tasks are mainly conducted manual or by automatic methods, such as the learning-from-crowds algorithms [108] to estimate the ground-truth labels from noisy crowd-sourced labels. To reduced experts' efforts or further improve the results of automatic methods, there are some works that employ visual analytics techniques to interactively improve the data quality. Tab. 1 shows that in recent years, this topic gains more and more research attention.

Feature engineering is used to select the best features to train the model. For example, in computer vision, we could use HOG (Histogram of Gradient) features instead of using raw image pixels. In visual analytics, interactive feature selection aims to form an interactive and iterative feature selection process. In recent years, in the deep learning era, feature selection and construction is mostly conducted via neural networks. Echoing this trend, there is less research attention in recent years (2016-2020) in this direction (Tab. 1).

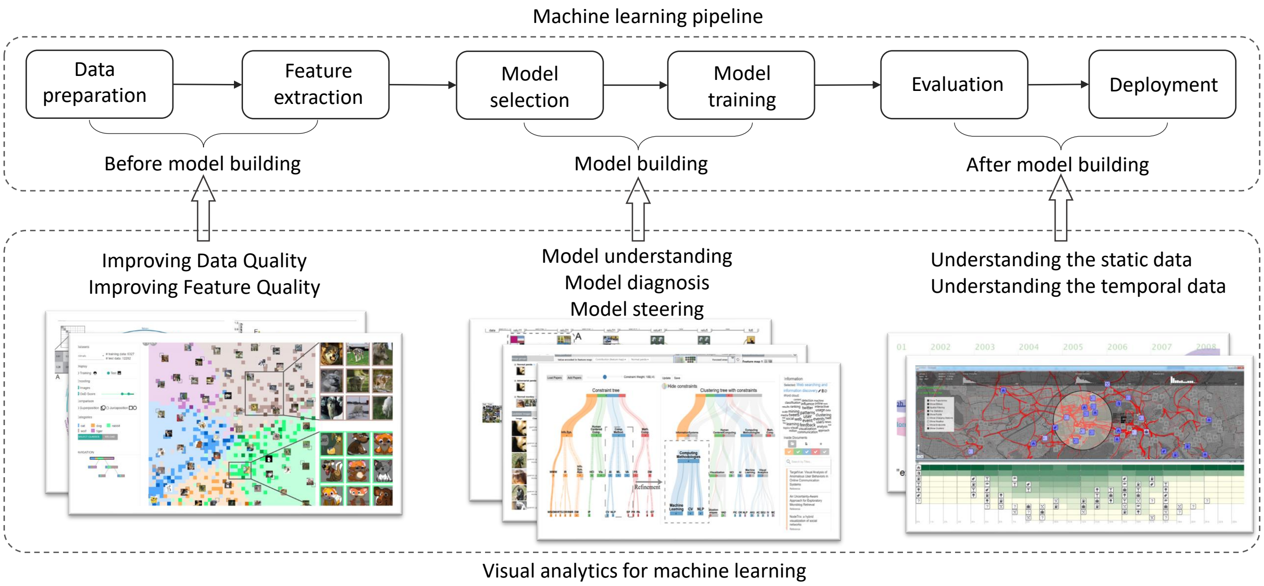

Fig. 1 An overview of visual analytics research for machine learning.

2.2.2 Techniques in Model Building

Model building is a central stage in building a successful machine learning application. Developing visual analytics methods to facilitate model building is also a growing research direction in visualization (Tab. 1). In this survey, we categorize current methods by their analysis goal, i.e., model understanding, diagnosis, and steering. Model understanding aims at visually explaining the working mechanisms of a model, such as how the change of parameters will influence the model and why the model gives an output for a specific input. Compared with model understanding, model diagnosis focuses on the modeling training process. It mainly targets at diagnosing the errors in training via interactive exploration of the training process. Model steering is mainly for interactively improving the model performance. For example, to refine a topic model, Utopian [53] enables users to interactively merge or split topics, and automatically modify the other topics accordingly.

2.2.3 Techniques after Model Building

After a machine learning model is built and deployed, it is crucial to help model users (e.g., a domain expert) to understand the model output in an intuitive way and further enable the trust of the model output. To this end, there are many visual analytics methods developed for different applications to explore the model output. Compared with the methods for model understanding in model building, these methods often target model users instead of model developers. Thus, the internal workings of a model are not illustrated, but the focus is on the intuitive presentation and exploration of model output. As these methods are often data-driven or application-driven, in this survey, we categorize these methods by the type of data being analyzed, i.e., static data or temporal data.

3 Techniques before Model Building

Two major tasks before building a model are data processing and feature engineering. They are critical because practical experience indicates that low-quality data and features will degrade the performance of machine learning models [196, 242]. Data quality issues include missing values, outliers, and noises in instances and their labels. Feature quality issues are represented as irrelevant features, redundancy among features, etc. While manually addressing these issues is time-consuming, automatic methods would also face poor performance. To this end, various visual analytics techniques have been developed to reduce experts' efforts as well as to simultaneously improve the performance of automatic methods for high-quality data and features.

3.1 Improving Data Quality

Data includes instances and their labels [198]. From this perspective, existing efforts for improving data quality can be classified into two categories: (1) instance-level improvement, and (2) label-level improvement.

Instance-Level Improvement. At the instance

Tab. 1 Categories of visual analytics techniques for machine learning and the representative works in each category.The number of papers is shown in the brackets.

Technique categories		**Papers**	**Trend**
Techniques before Model Building	Improving Data Quality (31)	[3], [11], [14], [16], [17], [18], [25], [45], [61], [91], [96], [101], [102], [118], [123], [125], [136], [144], [156], [192], [201], [203], [204], [213], [227], [228], [231], [256], [258], [267], [274]	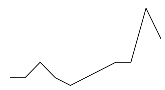
	Improving Feature Quality (6)	[109], [132], [183], [194], [222], [238]	
Techniques in Model Building	Model Understanding (30)	[28], [38], [56], [71], [79], [84], [104], [115], [116], [119], [120], [137], [155], [158], [184], [187], [188], [195], [197], [209], [212], [219], [223], [234], [252], [253], [254], [268], [289], [297]	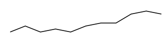
	Model Diagnosing (19)	[2], [7], [19], [32], [33], [63], [86], [98], [131], [154], [163], [177], [206], [210], [229], [233], [251], [266], [291]	
	Model Steering (29)	[23], [39], [40], [41], [53], [60], [64], [69], [73], [76], [77], [127], [138], [141], [153], [169], [180], [185], [189], [193], [200], [207], [216], [220], [245], [249], [261], [282], [296]	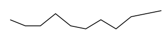
Techniques after Model Building	Understanding static data analysis results(43)	[4], [15], [22], [27], [29], [37], [43], [47], [57], [66], [72], [75], [81], [83], [85], [87], [89], [92], [100], [105], [106], [107], [112], [113], [114], [117], [121], [126], [128], [129], [146], [159], [160], [166], [199], [202], [205], [224], [247], [250], [276], [293], [295]	
	Understanding dynamic data analysis results (101)	[5], [6], [8], [9], [10], [20], [24], [26], [31], [34], [35], [36], [42], [46], [48], [49], [50], [51], [52], [55], [58], [59], [62], [67], [68], [70], [74], [78], [80], [88], [90], [93], [94], [95], [97], [99], [110], [111], [122], [124], [133], [134], [135], [140], [143], [145], [147], [150], [151], [164], [165], [167], [168], [173], [174], [175], [176], [179], [181], [191], [208], [211], [215], [218], [225], [230], [232], [235], [236], [237], [240], [241], [243], [246], [248], [259], [260], [263], [264], [265], [269], [270], [271], [272], [273], [275], [277], [278], [279], [280], [283], [284], [286], [287], [288], [290], [292], [294], [298], [299], [300]	

level, many visual analytics methods focus on detecting and correcting anomalies in data, such as missing values and duplication. For example, Kandel *et al.* [123] proposed Profiler to aid the discovery and assessment of anomalies in tabular data. Anomaly detection methods are applied to detect data anomalies, which are classified into different types subsequently. Based on the detected anomalies and their types, linked summary visualizations are automatically recommended to facilitate the discovery of potential causes and consequences of these anomalies. VIVID [3] was developed to handle missing values in longitudinal cohort study data. Through multiple coordinated visualizations, experts can identify the root cause of

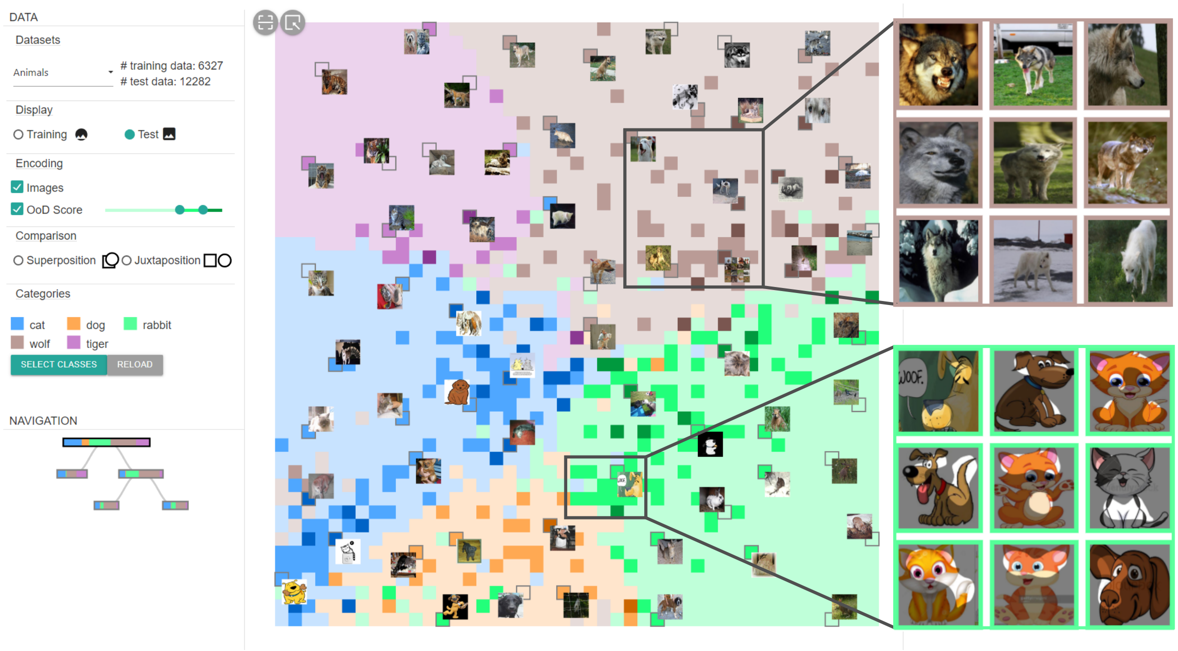

Fig. 2 OoDAnalyzer [45], an interactive method to detect out-of-distribution samples and explain them in context.

missing values (*e.g.*, a particular group who does not participate in follow-up examinations), and replace the missing data with an appropriate imputation model. Anomalies removal is often an iterative process and is performed multiple times. Illustrating the provenance in this iterative process allows users to be aware of the change of data quality and build trust in the processed data. To this end, Bors *et al.* [25] proposed DQProv Explorer to support the analysis of data processing provenance. It consists of a provenance graph to support the navigation of data states and a quality flow to present the change of data quality over time. Recently, another type of data anomaly, out-of-distribution (OoD) samples, has received extensive attention [139, 142]. OoD samples are test samples that are not well covered by training data, which is a major source of model performance degradation. To tackle this issue, Chen *et al.* [45] proposed OoDAnalyzer to detect and analyze OoD samples. An ensemble OoD detection method, combining both high- and low-level features, was proposed to improve the detection accuracy. Based on the detection result, a grid visualization (Fig. 2) is utilized to explore OoD samples in context and explain the underlying reason for their appearance. In order to generate grid layouts in interactive rates during the exploration, a kNN-based grid layout algorithm motivated by Hall's theorem was developed.

When considering time-series data, several challenges arise as time has distinct characteristics that induce specific quality issues and require analysis in a temporal context. To tackle this issue, Arbesser *et al.* [11] proposed a visual analytics system, Visplause, to visually assess time-series data qualities. Anomaly detection results, *e.g.*, frequencies of anomalies and their temporal distributions, are shown in a tabular layout. In order to address the scalability problem, data are aggregated in a hierarchy based on meta-information, which enables analyzing a group of anomalies (*e.g.*, abnormal time series of the same type) simultaneously. Besides automatically detected anomalies, KYE [91] also supports the identification of additional anomalies that have been overlooked by automatic methods. Time-series data is presented in a heatmap view, where abnormal patterns (*e.g.*, regions with unusually high values) are potential anomalies. Clickstream data is a widely studied time-series data in the field of visual analytics. To better analyze and refine clickstream data, Segmentifier [61] was proposed to provide an iterative exploration process for segmenting and analyzing. Users can explore segments in three coordinated views at different granularities and refine them by filtering, partitioning, and transforming. Every refinement step results in new segments, which can be further analyzed and refined.

To tackle the uncertainties in the data quality improvement, Bernard *et al.* [16] developed a visual analytics tool to exhibit the changes in the data and uncertainties caused by different preprocessing methods. With this tool, experts are aware of the effects of these methods and choose suitable ones to reduce task-irrelevant parts while preserving task-relevant parts of the data.

As data has the risk of exposing sensitive information, several recent studies have focused on preserving data privacy in the data quality improvement process. For tabular data, Wang *et al.* [258] developed a Privacy Exposure Risk Tree to display privacy exposure risks in the data and a Utility Preservation Degree Matrix to exhibit how the utility is changed as privacy-preserving operations are applied. To preserve privacy in network datasets, Wang *et al.* [256] presented a visual analytics system, *GraphProtector*. To preserve important structures of networks, node priorities are first specified based on their importance. Important nodes are assigned with low priorities, reducing the possibility of these nodes to be modified. Based on node priorities and utility metrics, users can apply and compare a set of privacy-preserving operations and choose the "best" one according to their knowledge and experiences.

Label-Level Improvement. According to whether the data has noisy labels, existing works can be classified into two categories: improving the quality of noisy labels and interactive labeling.

Crowdsourcing provides a cost-effective way to collect noisy labels. However, the annotations provided by crowd workers are usually noisy [152, 242]. Many methods have been proposed to remove noise in the labels. Willett *et al.* [267] developed a crowd-assisted clustering method to remove redundant explanations provided by crowd workers. Explanations are clustered into groups, of which the most representative ones are preserved. Park *et al.* [204] proposed C^{2}A that visualizes crowdsourced annotations and worker behaviors to help doctors identify malignant tumors in clinical videos. With C^{2}A, doctors can discard most tumor-free video segments and focus on the ones that are likely to contain tumors. To analyze the accuracy of crowdsourcing workers, Park *et al.* [203] developed CMed that visualizes clinical image annotations by crowdsourcing and workers' behaviors. By clustering workers according to their annotation accuracy and analyzing their logged events, experts are able to find good workers and observe the effect of workers' behavior patterns. LabelInspect [156] was proposed to improve crowdsourced labels by validating uncertain instance labels and unreliable workers. Three coordinated visualizations, a confusion (Fig. 3(a)), an instance (Fig. 3(b)), and a worker visualization (Fig. 3(c)), were developed to facilitate the identification and validation of uncertain instance labels and unreliable workers. Based on expert validation, more instances and workers are recommended for validating by an iterative and progressive verification procedure.

Although the aforementioned methods can effectively improve crowdsourced labels, such crowd information are not available in many real-world datasets. For example, the ImageNet dataset [214] only contains the cleaned labels via automatic noise removal methods. To tackle these datasets, Xiang *et al.* [274] developed DataDebugger to interactively improve data quality by utilizing user-selected trusted items. A hierarchical visualization combined with an incremental projection method and an outlier biased sampling method are presented to facilitate the exploration and identification of trusted items. Based on these identified trusted items, a data correction algorithm was developed to propagate labels from trusted items to the whole dataset. Paiva *et al.* [201] assumed that instances misclassified by a trained classifier were likely to be mislabeled instances. Based on this assumption, they employed a Neighbor Joining Tree enhanced by multidimensional projections to help users explore misclassified instances and correct mislabeled ones. After the correction, the classifier can be refined with the corrected labels, and a new round of correction starts. Bäuerle *et al.* [14] developed three classifier-guided measures to detect data errors. Data errors are then presented in a matrix and a scatter plot, which allows experts to reason about and resolve errors.

All the above methods start with a set of labeled data with noise. However, many datasets do not contain such a label set. To tackle this issue, many visual analytics methods have been proposed for interactive labeling. Reducing labeling efforts is one major goal of interactive labeling. To this end, Moehrmann *et al.* [192] used a SOM-based visualization to place similar images together, which allows users to label multiple similar images of the same class in one go. This strategy is also used by Khayat *et al.* [125] to identify social spambot groups with similar anomaly behaviors, Kurzhals *et al.* [136] to label mobile eye-tracking data, and Halter *et al.* [96] to annotate and analyze primary color strategies used in films. Apart from placing similar items together, other strategies, like filtering, are also applied to find items of interest for labeling. Filtering and sorting are utilized in MediaTable [213] to find similar video segments. A table visualization is

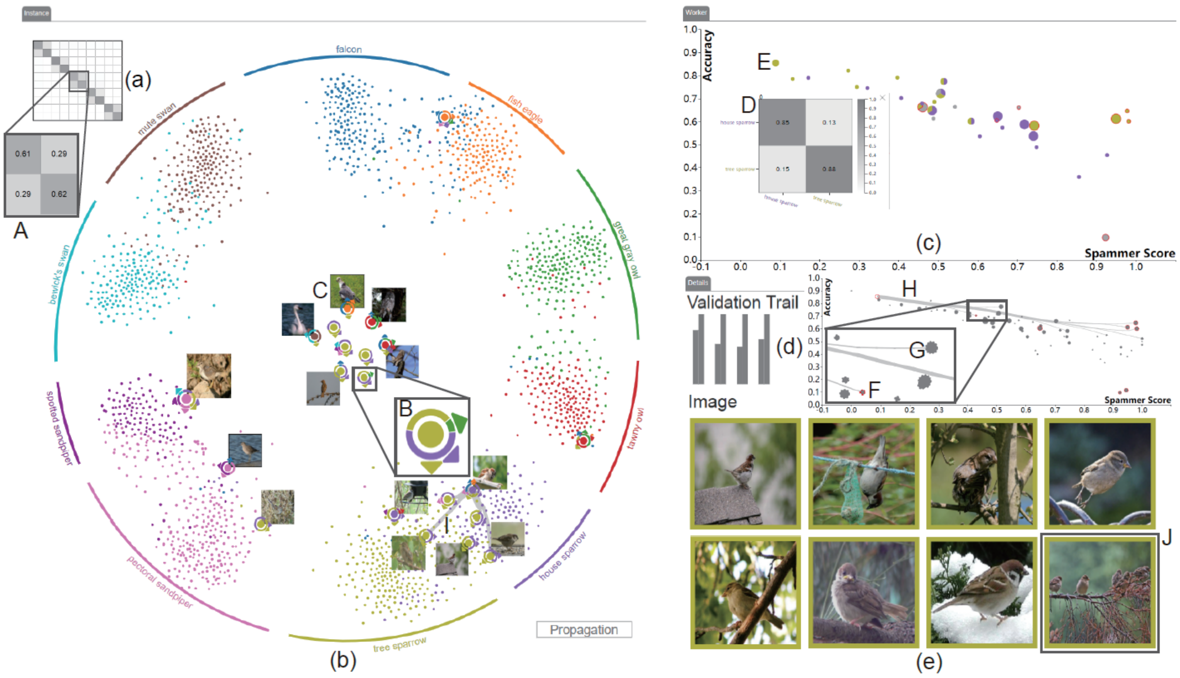

Fig. 3 LabelInspect [156], an interactive method to verify uncertain instance labels and unreliable workers.

utilized to present video segments and their attributes. Users can filter out irrelevant segments and sort on attributes to order relevant segments. Consequently, users are able to label several segments of the same class simultaneously. Stein *et al.* [231] provided a rule-based filtering engine to find the patterns of interest in soccer match videos. Experts can interactively specify rules through a natural language GUI.

Recently, to enhance the effectiveness of interactive labeling, various visual analytics methods combine visualization techniques with machine learning techniques, such as active learning. The concept of "intra-active labeling" was first introduced by Hoferlin *et al.* [102], which enhances active learning with human knowledge. Users are not only able to query instances and label them via active learning, but also to understand and steer machine learning models interactively. This concept is also used in text document retrieval [101], sequential data retrieval [144], trajectory classification [118], identifying relevant tweets [227], and argumentation mining [228]. For example, to annotate text fragments in argumentation mining tasks, Sperrle *et al.* [228] developed a language model for fragment recommendation. A layered visual abstraction is utilized to support five relevant analysis tasks required by text fragments annotation. In addition to developing systems for interactive labeling, some empirical experiments were conducted to demonstrate its effectiveness. For example, Bernard *et al.* [17] conducted experiments to show the superiority of user-centered visual interactive labeling over model-centered active learning. A quantitative analysis [18] was also performed to evaluate user strategies for selecting samples in the labeling process. Results show that in early phases, data-based user strategies (*e.g.*, clusters and dense areas) work well. However, in later phases, model-based user strategies (*e.g.*, class separation) perform better.

3.2 Improving Feature Quality

A typical method to improve feature quality is selecting useful features that contribute most to the prediction, i.e., feature selection [44]. A common feature selection strategy is to select a subset of features that minimizes the redundancy among them and maximizes the relevance between them and targets (*e.g.*, classes of instances) [183]. Along this line, several methods have been developed to interactively analyze the redundancy and relevance of features. For example, Seo *et al.* [222] proposed a rank-by-feature framework, which ranks features by relevance. They visualized ranking results with tables and matrices. Ingram *et al.* [109] proposed a visual analytics system, DimStiller, which allows users to explore the features and their relationships and interactively remove irrelevant and redundant features. May *et al.* [183] proposed SmartStripes to select different feature subsets for

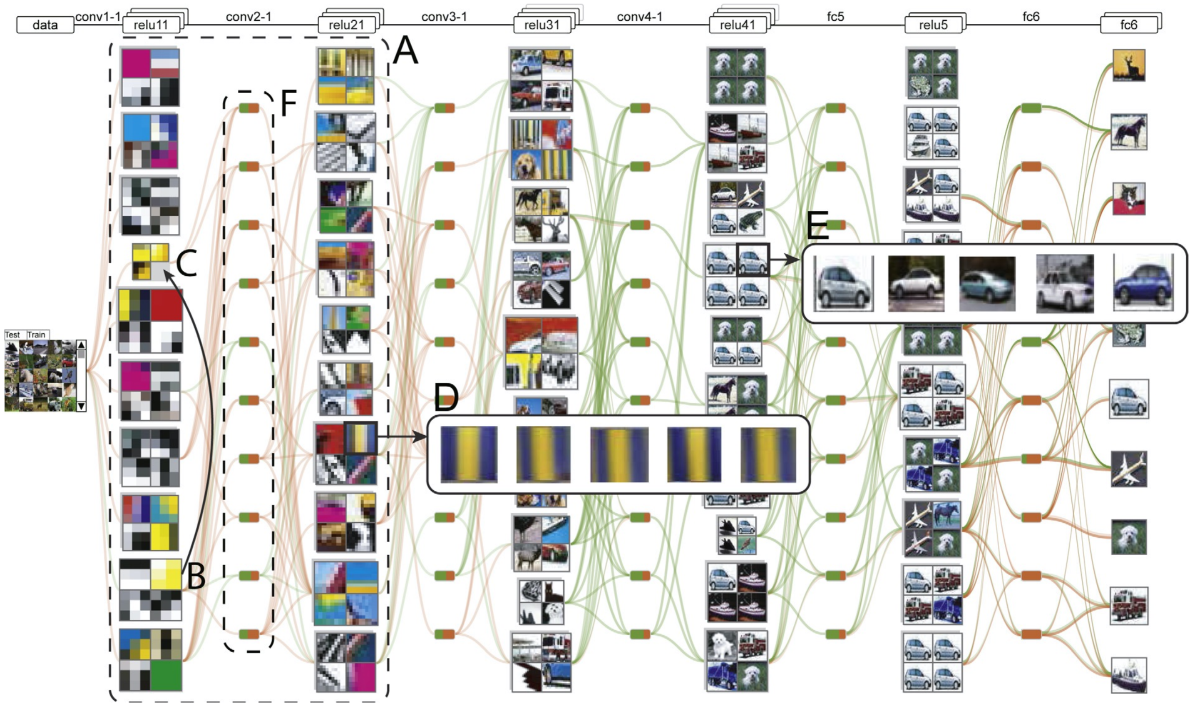

Fig. 4 CNNVis [155], a network-centric visual analytics technique to understand deep convolutional neural networks with millions of neurons and connections.

different data subsets. A matrix-based layout is utilized to exhibit the relevance and redundancy of features. Mühlbacher *et al.* [194] developed a partition-based visualization for the analysis of the relevance of features or feature pairs. The features or feature pairs are partitioned into subdivisions, which allows users to explore the relevance of features (or feature pairs) at different levels of details. A parallel coordinates visualization was utilized by Tam *et al.* [238] to identify features that could discriminate different clusters. Krause *et al.* [132] ranked features across different feature selection algorithms, cross-validation folds, and classification models. Users are able to interactively select the features and models that lead to the best performance.

Besides selecting existing features, constructing new features is also useful to facilitate the model building. For example, FeatureInsight [30] was proposed to construct new features for text classification. By visually examining the classifier errors and summarizing the root causes of these errors, users are able to create new features that can correctly discriminate these misclassified documents. To improve the generalization capability of new features, visual summaries are used to analyze a set of errors instead of individual ones.

4 Techniques in Model Building

Machine learning models are usually regarded as "black boxes" because of their lack of interpretability, which hinders their practical use in risky scenarios such as self-driving cars and financial investment. Current visual analytics techniques in model building explore to reveal the underlying working mechanisms of machine learning models and then help model developers to build well-performed models. First of all, model developers call for a comprehensive understanding of the models in order to release them from the time-consuming trial-and-error process. When the training process crashes or the model fails to get satisfying performance, model developers demand to diagnose the issues occurring in the training process. Finally, there is a need to assist model steering as much time is spent in improving the model performance in the model building process. Echoing these needs, researchers developed many visual analytics methods to enhance model understanding, diagnosing, and steering [54, 162].

4.1 Model Understanding

The works related to model understanding can be categorized into two classes: (1) understanding the effect of parameters, and (2) understanding the model behaviours.

Understanding the effect of parameters. One aspect of model understanding is to inspect how the model outputs change as the model parameters change. For example, Ferreira *et al.* [79] developed BirdVis to explore the relationships between different parameter configurations and model outputs, i.e., bird occurrence predictions in their application. The tool also reveals how these parameters are related to each other in the prediction model. Zhang *et al.* [292] proposed a visual analytics method to visualize how the variables affect the statistical indicators in the logistic regression model.

Understanding the model behaviours. Another aspect is to figure out how the model works to produce desired outputs. There are mainly three types of methods to explain the model behaviours, namely network-centric, instance-centric, and hybrid methods. Network-centric methods aim at exploring the model structure and interpreting how different parts of the model (*e.g.*, neurons/layers in convolutional neural networks) cooperate with each other to produce the final outputs. Earlier works employ directed graph layouts to visualize the structure of neural networks [244], but visual clutter remains a serious problem with the model structure becoming increasingly complex. To tackle this problem, Liu *et al.* [155] developed CNNVis to visualize deep convolutional neural networks. CNNVis (Fig. 4) leverages clustering techniques to group neurons of similar roles as well as their connections in order to address visual clutters caused by the huge amount.This tool helps experts understand the roles of the neurons and their learned features, and moreover, the process of how low-level features are aggregated into high-level ones through the network. Later, Wongsuphasawat *et al.* [268] designed a graph visualization for exploring the machine learning model architecture in Tensorflow [1]. They conducted a series of graph transformations to compute a legible interactive graph layout from a given low-level dataflow graph to display the high-level structure of the model.

Instance-centric methods aim at providing instance-level analysis and exploration, as well as understanding the relationships among instances. Rauber *et al.* [209] visualized the representations learned from each layer in the neural network by projecting them onto 2D scatterplots. Users can identify clusters and confusion areas in the representation projections and, therefore, understand the representation space learned by the network. Furthermore, they can study how the representation space is evolved through the training session so as to understand the network's learning behaviour. Some visual analytics techniques for understanding recurrent neural networks (RNNs) also adopts such a instance-centric design. LSTMVis [234] developed by Strobelt *et al.* utilizes parallel coordinates to present the hidden states, which supports the analysis of the changes in the hidden states over texts. RNNVis [187] developed by Ming *et al.* clusters the hidden state units (a hidden state unit refers to one dimension of the hidden states vector in RNNs) as memory chips and words as word clouds. Their relationships are modeled as a bipartite graph, which supports the sentence-level explanations in RNNs.

Hybrid methods combine the above two methods and leverage both of their strengths. In particular, the instance-level analysis can be enhanced with the context of the network architecture. Such contexts will benefit the understanding of the network's working mechanism. For instance, Hohman *et al.* [104] proposed Summit to reveal important neurons and critical neuron associations contributing to the model prediction. It integrates an embedding view to summarize the activations among classes and an attribute graph view to reveal the influential connections between neurons. Kahng *et al.* [119] proposed ActiVis to large-scale deep neural networks. It visualizes the model structure with a computational graph and the activation relationships among instances, subsets, and classes using a projected view.

In recent years, there have been some efforts to use a surrogate explainable model to explain the model behaviour. The major benefit of these methods is that they do not require users to get into the model itself. Thus, they are more applicable to those without or with limited machine learning knowledge. Treating the classifier as a black box, Ming *et al.* [188] extracted the rule-based knowledge from the input and output of the classifier first. These rules are then visualized using RuleMatrix, which supports practitioners to interactively explore the extracted rules and improve the interpretability of the model. Wang *et al.* [253] developed DeepVID to generate the visual interpretation for image classifiers. Given an image of interest, a deep generative model was first used to generate the samples near it. These generated samples were used to train a simpler and more interpretable model, such as a linear regression classifier, which helps explain how the original model makes the decision.

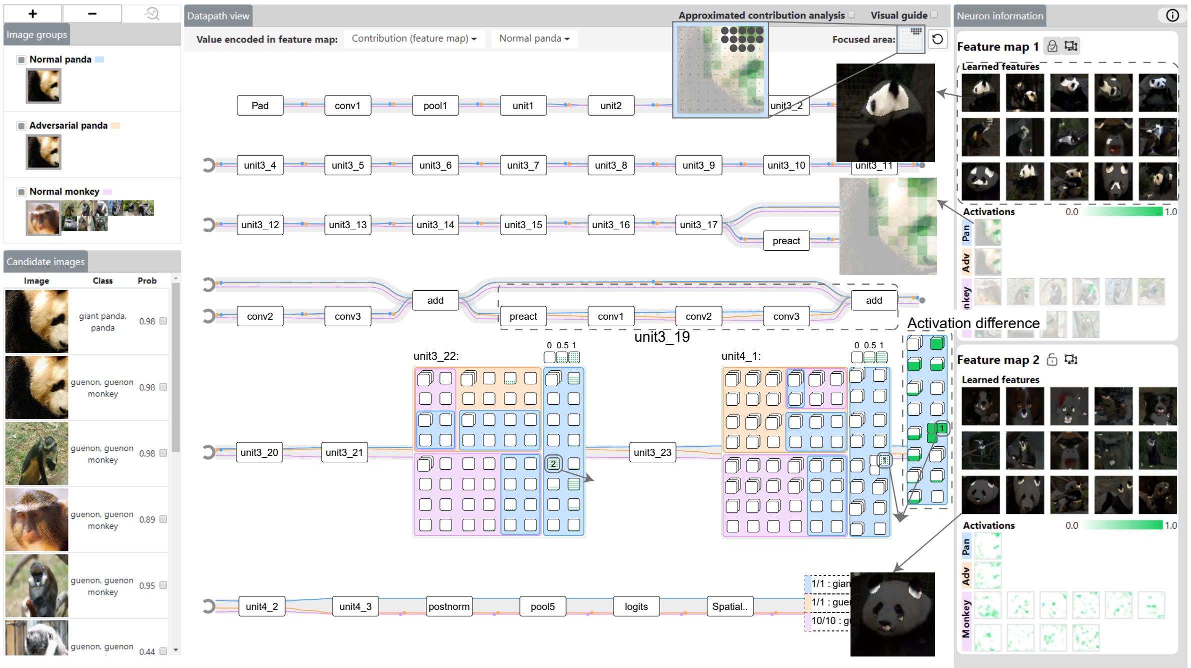

Fig. 5 AEVis [33], a visual analytics system for analyzing adversarial samples. It shows the diverging and merging patterns in the extracted datapaths with a river-based visualization and the critical feature maps with a layer-level visualization.

4.2 Model Diagnosing

According to the content involved in the diagnosis, existing visual analytic techniques for model diagnosing fall into two categories: (1) analyzing the training results, and (2) analyzing the training dynamics.

Analyzing the training results. There are tools developed for the diagnosis for classifiers based on the unsatisfying performance [7, 19, 86, 210]. For example, Squares [210] used boxes to represent samples and group them according to their prediction classes. With different textures that encode true/false positives/negatives, this tool is proved to allow fast and accurate estimation of performance metrics at multiple levels of detail. Recently, the issue of model fairness has drawn growing attentions [2, 32, 266]. For example, Ahn *et al.* [2] proposed a framework named *FairSight* and implemented a visual analytics system to support the analysis of fairness in ranking problems. They divided the machine learning pipeline into three phases (data, model, and outcome) and then measured the bias both at the individual level and group level using different measures. Based on these measures, developers can iteratively identify the features that cause discrimination and remove them from the model. Researchers are also interested in exploring the potential vulnerabilities in the model that prevent it from being reliably applied to real-world applications [33, 177]. Cao *et al.* [33] proposed AEVis to analyze how the adversarial examples fooled the neural networks. The system (see Fig. 5) takes both normal and adversarial examples as the input and extracted their datapaths for the model prediction. It then employs a river-based metaphor to show the diverging and merging patterns of the extracted datapaths, which reveals where the adversarial samples mislead the model. Ma *et al.* [177] designed a series of visual representations to reveal how data poisoning will make a model misclassify a specific sample from overview to detail. By comparing the distribution of the poisoned and the normal training data, experts can conclude the reason for the misclassification of the attacked sample.

Analyzing the training dynamics. Recent efforts have also been concentrated on analyzing the training dynamics. These techniques are developed for debugging the training process of machine learning models. For example, DGMTracker [154] assists experts to reason the failed training process of deep generative models. It utilizes a blue-noise polyline sampling algorithm to simultaneously keep the outliers and the major distribution of the training dynamics in order to help experts detect the potential root cause of a failure. It also employs a credit assignment algorithm to disclose the interaction among neurons to facilitate the diagnosis of the propagation of the

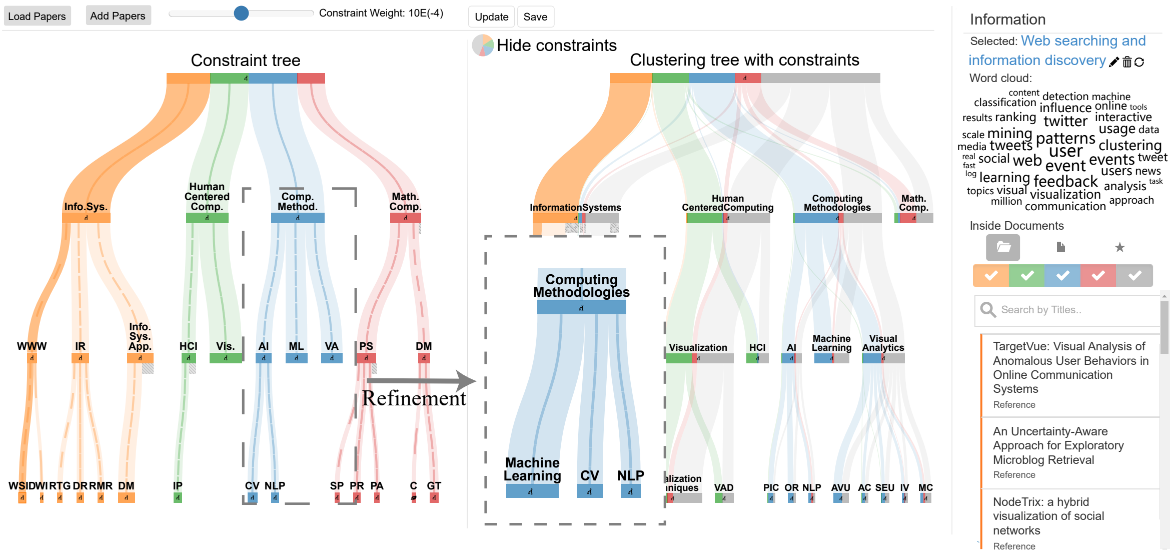

Fig. 6 ReVision [282], a visual analytics system integrating a constrained hierarchical clustering algorithm with an uncertainty-aware, tree-based visualization to help users interactively refine the hierarchical topic modeling results.

failure. Besides, attention has been brought to the diagnosis of the training process of deep reinforcement learning. Wang *et al.* [251] proposed DQNViz for the understanding and diagnosis of deep Q-networks for a Breakout game. At the overview level, DQNViz presents the changes in the overall statistics over the training process with line charts and stacked area charts. Then at the detail level, it uses segment clustering and then a pattern mining algorithm to help experts identify the common as well as the suspicious patterns in the event-sequences of the agents in Q-networks. As another example, He *et al.* [98] proposed DynamicsExplorer to diagnose LSTM trained for controlling a ball-in-maze game. To support quick identification of where the training failure arises, it visualizes the ball trajectories with a trajectory variability plot as well as their clusters with a parallel coordinates plot.

4.3 Model Steering

There are two major strategies for model steering: (1) refining the model with human knowledge, and (2) selecting the best model from a model ensemble.

Refining the model with human knowledge. Several visual analytics techniques have been developed to loop users into the model refinement process through flexible interactions.

Users can directly refine the target model with visual analytics techniques. A typical example is ProtoSteer [189], a visual analytics system that enabled editing prototypes to refine a prototype sequence network named ProSeNet [190]. ProtoSeer uses four coordinated views to present the information of the learned prototypes in ProSeNet. Users can refine these prototypes by adding, deleting, and revising some specific prototypes. The model will then be retrained with these user-specific prototypes for performance gain. In addition, van der Elzen *et al.* [245] proposed BaobabView to support experts to construct decision trees iteratively with domain knowledge. Experts can refine the decision tree with direct operations, including growing, pruning, and optimizing the internal nodes, and evaluate the refined one with various visual representations.

Besides the direct model update, users can also correct the flaws in the results or provide extra knowledge, and the model will be updated implicitly to produce improved results based on human feedback. Several works have focused on incorporating user knowledge into topic models to improve their results [53, 69, 73, 127, 261, 282]. For instance, Yang *et al.* [282] presented ReVision that allows users to steer the hierarchical clustering results by leveraging an evolutionary Bayesian rose tree clustering algorithm with constraints. As shown in Fig. 6, the constraints and the clustering results are displayed with an uncertainty-aware tree-based visualization to guide the steering of the clustering results. Users can refine the constraint hierarchy by dragging. The documents will be re-clustered based on the modified constraints. Other human-in-the-loop models have also stimulated the development of visual analytic systems to support

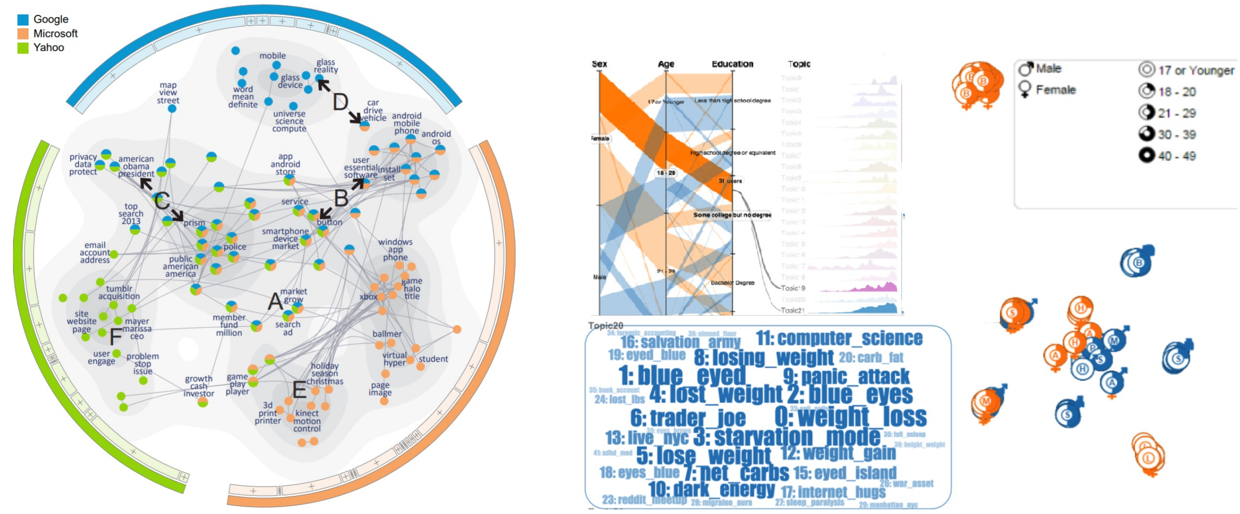

(a) (b)

Fig. 7 Some examples of static text visualization. (a) TopicPanorama [160] extracted the topic graphs from multiple sources and revealed the relationships among them using graph layout. (b) DemographicVis [66] measured the similarity between different users after analyzing their post contents, and revealed their relationships using t-SNE projection.

such kind of model refinement. For instance, Liu *et al.* [153] proposed MutualRanker using an uncertainty-based mutual reinforcement graph model to retrieve important blogs, users, and hashtags from microblog data. It shows the ranking results, uncertainty, and its propagation with the help of a composite visualization, and users can examine the most uncertain items in the graph and adjust their ranking scores. The model will incrementally update by propagating the adjustments through the graph.

Selecting the best model from a model ensemble. Another strategy for model steering is to select the best one from a model ensemble, which is usually found in clustering [41, 200, 220] and regression models [23, 60, 169, 207]. Clustrophile 2 [41] is a visual analytics system for visual clustering analysis, which guides users to select the appropriate input features and clustering parameters through recommendations based on user-interested results. BEAMES [60] was designed for multimodel steering and selection in regression tasks. It created a collection of regression models by varying algorithms and their corresponding hyperparameters, with further optimization by the interactive weighting of data instances and interactive feature selection and weighting. Users can inspect them and then select an optimal model according to different aspects of performance, such as their residual scores and mean squared errors.

5 Techniques after Model Building

Existing visual analytics efforts after model building aim at helping users understand and gain insights from model outputs, such as high-dimensional data analysis results [157, 161]. As these methods are often data-driven, we categorize the corresponding methods by the type of data to be analyzed. The temporal property of data is critical in visual design. Thus, in this paper, we classify the data into two categories: understanding static data analysis results and understanding dynamic data analysis results. A visual analytics system for understanding static data analysis results usually treats all the model output as a large collection and mainly analyzes the static structure. For dynamic data, in addition to understanding the analysis results at each time point, the system focuses more on illustrating the evolution of data over time, which is learned by the analysis model.

5.1 Understanding Static Data Analysis Results

We summarize the research on understanding static data analysis according to data types. Most research efforts focus on textual data analysis, while fewer works study the understanding of other types of data analysis.

Textual data analysis. The most widely studied topic is visual text analytics, which tightly integrates interactive visualization techniques with text mining techniques (e.g., document clustering, topic models, and word embedding) to help users better understand and consume a large amount of textual data [161].

Some early works employed simple visualizations to directly convey the results of classical text mining techniques, such as text summarization, categorization, and clustering. For example, Görg *et al.* [89] developed a multi-view visualization consisting of a list view, a cluster view, a word cloud, a grid view, and a document view, to visually illustrate the analysis results of document summarization, document clustering, sentiment analysis, entity identification, and recommendation. By combining interactive visualization with text mining techniques, a smooth and informative exploration environment is provided to users.

Most of the later research has focused on combining well-designed interactive visualization with state-of-the-art text mining techniques, such as topic models and deep learning models, to disclose more insights embedded in textual data. To provide an overview of the relevant topics discussed in multiple sources, Liu *et al.* [160] first utilized the correlated topic model to extract the topic graphs from multiple text sources, respectively. A graph matching algorithm is then developed to match the topic graphs from different sources, and a hierarchical clustering method is employed to generate the hierarchies for topic graphs. Both the matched topic graph and hierarchies are fed into a hybrid visualization consists of a radial icicle plot and a density-based node-link diagram (Fig. 7(a)), to support the exploration and analysis of the common and distinctive topics discussed in multiple sources. Dou *et al.* [66] introduced DemographicVis to analyze different demographic groups on social media based on the content generated by users. An advanced topic model, latent Dirichlet allocation (LDA) [178], was employed to extract the topic features from the corpus. The relationships between the demographic information and extracted features were explored through a Parallel Sets visualization [130], and different demographic groups are projected onto the two-dimension space based on their similarity of the topics of interest (Fig. 7(b)). Recently, some deep learning models are also adopted because of better performance. For example, Berger *et al.* [15] proposed cite2vec to visualize the latent themes in a document collection via the document usage (*e.g.*, citation). It extended a famous word2vec model, skip-gram model [186], to generate the embedding for both words and documents by considering the citation information and the text content together. The words are projected onto two-dimension space using t-SNE first, and the documents are projected onto the same space, where both the document-word relationship and document-document relationships are considered simultaneously.

Other types of data analysis. In addition to textual data, other types of data are also studied. For example, Hong *et al.* [105] analyzed the flow field through an LDA model by defining pathlines as documents and features as words, respectively. After the modeling, the original pathlines and the extracted topics were projected onto a two-dimension space using multidimensional scaling, and several previews were generated to show the rendering result of the pathlines in the important topics. Recently, a visual analytics tool, SMARTexplore [22], was developed to help analysts find and understand interesting patterns within and across dimensions, including correlations, clusters, and outliers. To this end, it tightly couples a table-based visualization with pattern matching and subspace analysis.

5.2 Understanding Dynamic Data Analysis Results

In addition to understanding the results of static data analysis, it is also important to investigate and analyze how the latent themes in data change over time. For example, it is very helpful for politicians to make the decisions timely if the system provides an overview of the major public opinions on social media and how they change over time. Most of the existing works focus on understanding the analysis results of a data corpus where each data item is associated with a time stamp. According to whether the system supports the analysis of streaming data, we further classified the existing works on visual dynamic data analysis into two categories: offline analysis and online analysis. In offline analysis, all the data are available before analysis, while online analysis tackles streaming data that keeps on coming in the analysis process.

Offline analysis. Offline analysis research can be classified into three categories based on the analysis tasks: topic analysis, events analysis, and trajectory analysis.

Understanding the topic evolution of large text corpus over time is an important topic and attracts much attention. Most of the existing works adopted a river metaphor to convey the change of the text corpus over time. ThemeRiver [97] is one of the pioneer works, which uses the river metaphor to reveal the changes in the volume of different themes. To better understand the content change of a document corpus, TIARA [165, 264] utilizes an LDA model [21] to extract the topics from the corpus and reveal their change over time.

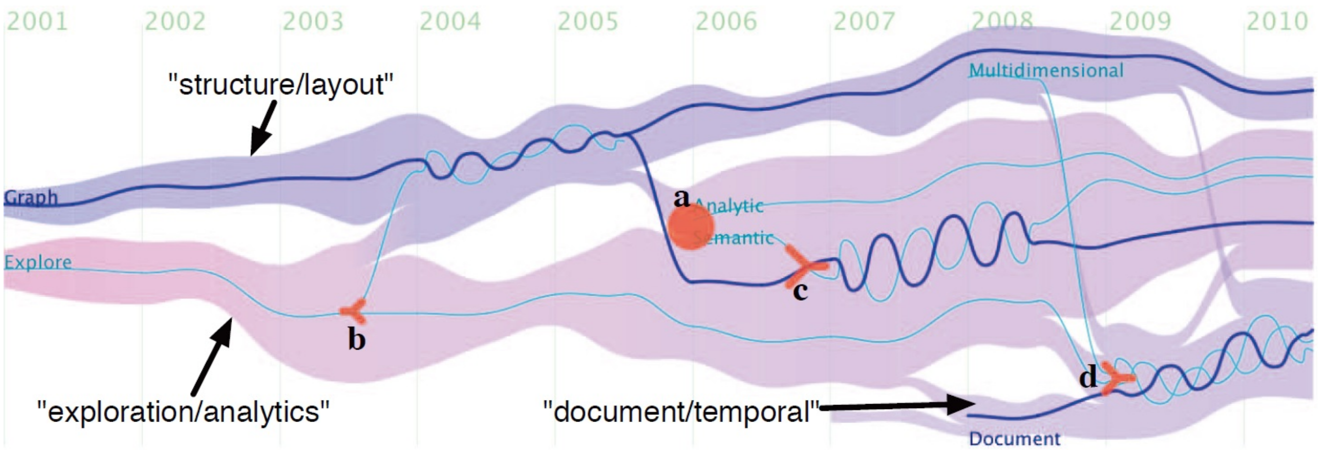

Fig. 8 TextFlow [58] employs a river-based metaphor to show topic birth, death, merging, and splitting.

However, only observing the volume and content change is not enough for the complex analysis tasks where users want to explore the relationship among different topics and their change over time. Therefore, later works have focused on understanding the relationships (e.g., topic splitting and merging) among topics and their evolving patterns along time. For example, Cui *et al.* [58] first extracted the topic splitting and merging patterns from a document collection by an incremental hierarchical Dirichlet process model [239]. Then a river metaphor with a set of well-designed glyphs was developed to visually illustrate the aforementioned topic relationships and their dynamic changes over time. Xu *et al.* [280] leveraged a topic competition model to extract the dynamic competition among topics and the effects of opinion leaders on social media. Sum *et al.* [236] extended the competition model to the coopetition (cooperation and competition) model to help understand the more complex interactions among the evolving topics. Wang *et al.* [260] proposed IdeaFlow, a visual analytics system for learning the lead-lag relationships across different social groups over time. However, the aforementioned works use a flat structure to model the topics, which hampers their usage in the era of big data for handling large-scale text corpora. Fortunately, there are already initial efforts in coupling hierarchical topic models with interactive visualization to favor the understanding of the main content in a large text corpus. For example, Cui *et al.* [59] extracted a sequence of topic trees using an evolutionary Bayesian rose tree algorithm [262] and then calculate the tree cut for each tree. These tree cuts are used as the approximations for the topic trees and displayed in a river metaphor, which also reveals the dynamic relationships among the topics, including topic birth, death, splitting and merging.

Event analysis targets at revealing common or semantically important sequential patterns in event sequence, which are ordered series of events [94, 112, 168, 176]. To facilitate visual exploration of large scale event sequence and pattern discovery, several visual analytics methods have been proposed. For example, Liu *et al.* [168] developed a visual analytics method for clickstream data. Maximal sequential patterns are discovered and pruned from the clickstream data. The extracted patterns and original data are well illustrated at four granularities: patterns, segments, sequences, and events. Guo *et al.* [94] developed EventThread, which uses a tensor-based model to transform the event sequence data into an n-dimension tensor. The latent patterns (threads) are extracted with a tensor decomposition technique, segmented into stages, and then clustered. These threads are represented as segmented linear stripes, and a line map metaphor is used to reveal the changes between different stages. Later, EventThread was extended to overcome the limitation of the fixed length of each stage [93]. The authors proposed an unsupervised stage analysis algorithm to effectively identify the latent stages in event sequences. Based on this algorithm, an interactive visualization tool is developed to reveal and analyze the evolution patterns across stages.

There are also some works focusing on understanding movement data (*e.g.*, GPS records) analysis results. Andrienko *et al.* [10] extracted the movement events from trajectories and then performed spatial-temporal clustering for aggregation. These clusters are visualized by the spatio-temporal envelopes to help analysts find the potential traffic jams in the city. Chu *et al.* [55] adopted an LDA model for mining the latent movement

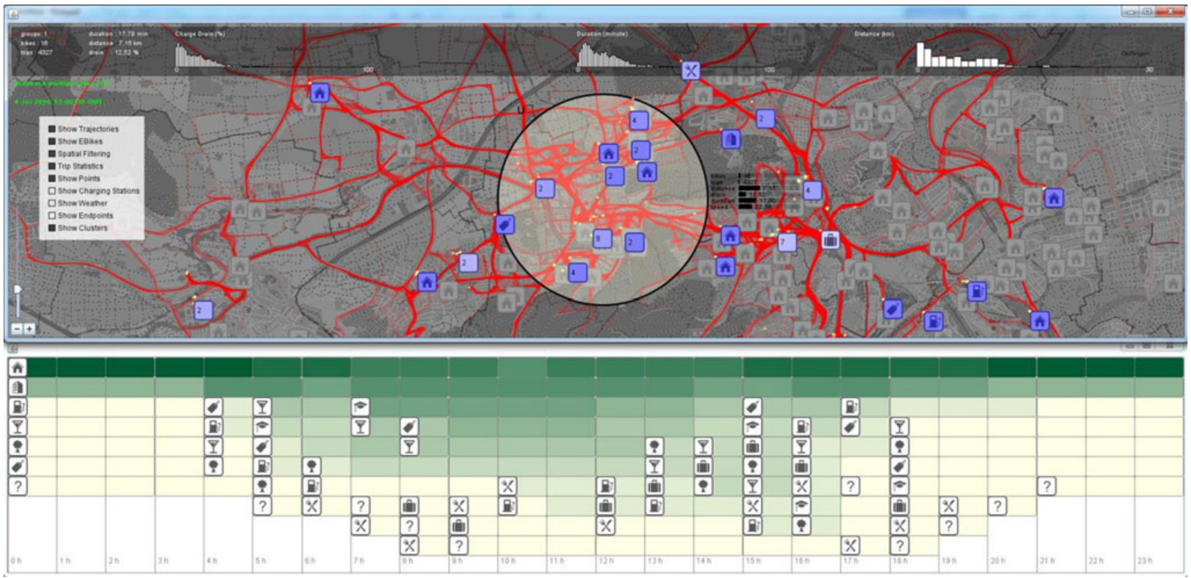

Fig. 9 Kruger *et al.* [135] enriched the trajectory data semantically. The frequent routes and destinations are visualized in the geographic view (top), while the frequent temporal patterns are mined and displayed in the temporal view (bottom).

patterns in taxi trajectories. The movement of each taxi represented by the traversed street names was regarded as a document. The parallel coordinate was used to visualize the distribution of streets over topics, where each axis represents a topic, and each polyline represents a street. The evolution of the topics was visualized as topic routes that connect similar topics between adjacent time windows. More recently, Zhou *et al.* [299] treated the origin-destination flows as words and trajectories as paragraphs, respectively. Therefore, a word2Vec model was used to generate the vectorized representation for each origin-destination flow. t-SNE was then employed to project the embedding of the flows onto two-dimension space, where analysts can check the distributions of the origin-destination flows and select some of them to be displayed on the map. Besides directly analyzing the original trajectories data, some other papers try to augment the trajectories with auxiliary information to reduce the burden on visual explorations. Kruger *et al.* [135] clustered the destinations with DBScan and then used Foursquare to provide detailed information about the destinations (*e.g.*, shops, university, residence). Based on the enriched data, the frequent patterns were extracted and displayed in the visualization (Fig. 9), where the icons on the time axis helped understand these patterns. Chen *et al.* [50] mined the trajectories from the geo-tagged social media and displayed the keywords extracted from the text content, which helps users explore the semantics of trajectories.

Online analysis. Online analysis is especially necessary for streaming data, such as text streams. As a pioneering work for online analysis of text streams, Thom *et al.* [240] proposed ScatterBlog to analyze geo-located tweet streams. The system uses Twitter4J to get streaming tweets and extract location, time, user ID, and tokenized terms in the tweets. To efficiently analyze a tweet stream, an incremental clustering algorithm was employed to cluster similar tweets. Based on the clustering results, spatial-temporal anomalies were detected and reported to the users in real-time. To reduce user efforts for filtering and monitoring in ScatterBlogs, Bosch *et al.* [26] proposed ScatterBlogs2, which enhanced ScatterBlogs with machine learning techniques. In particular, an SVM-based classifier was built for filtering tweets of interest, and an LDA model was employed to generate a topic overview. To efficiently handle high-volume text streams, Liu *et al.* [164] developed TopicStream to help users analyze hierarchical topic evolution in high-volume text streams. In TopicStream, an evolutionary topic tree was built from text streams, and a tree cut algorithm was developed to reduce visual clutter and enable users to focus more on topics of interest. Combining a river metaphor and a visual sedimentation metaphor, the tool effectively illustrates the overall hierarchical topic evolution as well as how new arriving textual documents are gradually aggregated into the existing topics over time. Triggered by TopicStream, Wu *et al.* [271] developed StreamExplorer, which enables the tracking and comparison of a social stream. In particular, an entropy-based event detection method was developed to detect the events from the social media stream. The events are further visualized in a multi-level visualization, including a glyph-based timeline, a map visualization, and interactive lenses. In addition to text streams, other types of streaming data are also analyzed. For example, Lee *et al.* [140] employed a long short-term memory model for road traffic congestion forecasting and visualized the results with a Volume-Speed Rivers visualization. The propagation of the congestions was also extracted and visualized, which helps analysts understand the causality within the detected congestions.

6 Research Opportunities

Although visual analytics research for machine learning has achieved promising results in both academic research and real-world applications, there are still several long-term research challenges. Here, we discuss and highlight major challenges and potential research opportunities in this area.

6.1 Opportunities before Model Building

Improving data quality for weakly supervised learning. Weakly supervised learning builds models from data with quality issues, including inaccurate labels, incomplete labels, and inexact labels. Improving data quality can boost the performance of weakly supervised learning models [148]. Most of the existing methods focus on inaccurate (*e.g.*, noisy crowdsourced annotations and label errors) data quality issues and interactive labeling related to incomplete (*e.g.*, none or only a few data are labeled) data quality issues. However, fewer efforts are devoted to the better exploitation of unlabeled data related to incomplete data quality issues as well as inexact (*e.g.*, coarse-grained labels that are not exact as required) data quality issues. This paves the way for potential future research.

First, the potential of visual analytics techniques to address the incomplete issue is not fully exploited. For example, improving the quality of unlabeled data is critical for semi-supervised learning [148, 149], which is tightly combined with a small amount of labeled data during the training to infer the correct mapping from the data set to the label set. One typical example is graph-based semi-supervised learning [149], which depends on the relationship between labeled and unlabeled data. Automatically constructed relationships (graphs) are sometimes poor in quality, resulting in model performance degradation. A major cause behind the poor-quality graphs is that automatic graph construction methods usually rely on global parameters (*e.g.*, the global k value of kNN graph construction method), which may not be appropriate for local regions. As a consequence, it is necessary to utilize visualization to illustrate how labels are propagated along the graph edges, thus facilitate the understanding of how local graph structures affect the model performance. Based on the understanding, experts can adaptively modify the graph and gradually create a higher-quality graph.

Second, although the inexact data quality issue is common in real-world applications [302], it receives less attention in the field of visual analytics. The inexact data quality issue refers to the situation where labels are inexact, *e.g.*, coarse-grained labels. One typical example of the inexact data quality issue is the coarse-grained labels of computed tomography (CT) scans. The labels of CT scans usually come from the corresponding diagnosis reports that describe whether patients have certain diseases (e.g., tumor). For a CT scan with tumors, we only know one or more slices in this scan contain tumors. However, we do not know which slices contain tumors as well as the exact locations on these slices. Although various machine learning methods [82, 301] have been proposed to learn from such coarse-grained labels, they may lead to poor performance [148] due to the lack of exact information. Fine-grained validations are still required to improve data quality. To this end, one potential solution is to combine interactive visualization with learning algorithms to better illustrate the root cause of bad performance by examining the overall data distribution and the wrongly predicted ones, and develop an interactive verification process for providing more fine-grained labels while minimizing expert efforts.

Explainable feature engineering. Most of the existing work for improving feature quality focus on tabular or textual data from traditional analysis models. The features of these data are naturally interpretable, which makes the feature engineering process much easier. In addition to these traditional feature representations, features extracted by deep neural networks perform better than the handcraft ones [65, 255]. However, these deep features are hard to interpret due to the "black box" nature of deep neural networks. The uninterpretable property of such features brings several challenges for feature engineering.

First, the extracted features are obtained in a data-driven process, which may not well represent the original images/videos when the datasets are biased. For example, given a dataset with only dark-colored dogs and light-colored cats, the extracted features only emphasize colors and ignore other discriminate concepts, like faces and ears. Without a clear understanding of these biased features, it is hard to correct them in a comprehensive way. Thus, an interesting topic for future work is to utilize interactive visualization to disclose why the features are biased. The key challenge here is how to measure the information preserved or abandoned by the extracted features and visualize them in a comprehensive manner.

Moreover, redundancy exists in extracted deep features [12]. Removing redundant features can also lead to several benefits, like reducing the storage requirement and improving generalization [44]. However, without a clear understanding of the exact meanings of features, it is hard to judge whether a feature is redundant. As a consequence, one interesting future work is to develop a visual analytics method to convey feature redundancy in a comprehensive way, allow experts to explore them, and remove redundant ones for better qualities.

6.2 Opportunities in Model Building

Online training diagnosis. Existing visual analytics tools for model diagnosing are mostly in an offline manner that the data for diagnosis is collected after the training process is finished. They have shown their capability for revealing the root causes of the failed training process. However, as modern machine learning models are becoming more and more complex, their training process can last for a few days or even several weeks. Such an offline manner severely restricts the efficiency of visual analytics to assist model diagnosis. Given this fact, it is of significant need to develop visual analytics tools to diagnose the online training process so that model developers can identify the anomalous training process and make corresponding adjustments to the potential issues promptly. This can save much time in the trial-and-error model building process. The key challenge for online diagnoses is to detect anomalies in the training process in time. While it remains a difficult task to develop algorithms for detecting anomalies automatically and accurately in the real-time environment, it will be applicable to leverage interactive visualization to locate the potential errors in the training process. Different from offline diagnoses, the data of the training process will be continuously fed into the online analysis tool. Thus, progressive visualization techniques are needed to produce meaningful visualization results of the partial streaming data. These techniques can help experts monitor the online model training process and identify possible issues rapidly.

Interactive model refinement. Recent works have explored the utilization of uncertainty to facilitate interactive model refinement [73, 153, 261, 282]. There are many methods to assign uncertainty scores to the model outputs (*e.g.*, based on the confidence scores produced by the classifiers), and then different visual hints can be used to guide users to examine the model outputs with high uncertainty. Models will re-compute the uncertainty after updated with the user refinement, and users can perform the refinement iteratively until they are satisfied with the results. Furthermore, additional information can also be leveraged in order to provide users with more intelligent guidance for achieving a fast and accurate model refinement process. However, the room for improving interactive model refinement is still largely unexplored for researchers. A possible direction is that since the refinement process usually requires several iterations, the guidance in the later iterations can be learned from users' previous interactions (*e.g.*, recommending new guidance learned from a series of previous user interactions). For example, in a clustering application, users may define some "must-link" or "cannot-link" constraints on some instance pairs, and such constraints can be used to instruct a model to split or merge some clusters in the intermediate result. In addition, it can be considered to predict where refinements will be needed based on some priori knowledge. For example, there might be some flaws if the model outputs conflict with some public or domain knowledge, especially in some unsupervised models (*e.g.*, Nonlinear Matrix Factorization and Latent Dirichlet Allocation for topic modeling), which should be noticed in the refinement process. Therefore, such a knowledge-based strategy focuses on revealing unreasonable results produced by the models, and then users can refine the models by adding constraints to the models.

6.3 Opportunities after Model Building

Understanding multi-modal data. Existing work on content analysis has achieved great success in understanding single-modal data, such as texts, images, and videos. However, real-world applications often contain multi-modal data, which is a combination of several different content forms, such as text, audio, and images.For example, in the medical scenario, a physician diagnoses the patient after synthetically analyzing multiple kinds of data, such as the medical record (text), laboratory report (tabular), and CT scanning (image). When analyzing these multi-modal data, the in-depth relationships among different modals can not be well captured by simply combining the learned knowledge from single-modal models. It is more promising to employ multi-modal machine learning techniques and leverage its capability to disclose insights across different forms of data. To this end, a more powerful visual analytics system is crucial for understanding the output of these multi-modal learning models.Many machine learning models are proposed to learn the joint representation of multi-modal data, including natural language, visual signals, and vocal signals [13, 170]. Accordingly, an interesting future direction is how to effectively visualize the learned joint representations among multi-modal in an all-in-one manner data. An effective visualization will facilitate the understanding of the multi-modal data and their relationships. Some classic multi-modal tasks can also be employed to enhance natural interactions in the field of visual analytics. For example, in the vision-and-language scenario, the visual grounding task (identify the corresponding image area given the description) can be used to provide a natural interface to support natural-language-based image retrieval in a visual environment.

Analyzing concept drift for better performance. In real-world applications, it is often assumed that the model mapping function from input data to output values (e.g., prediction label) is static. However, as data continues to come, the mapping between the input data and output values may change in unexpected ways [171]. Under such a situation, a model trained on historical data may no longer work properly on new data. This usually causes noticeable performance degradation for the application data that does not match the training data. Such a non-stationary learning problem over time is known as concept drift in the literature. As more and more machine learning applications directly consume streaming data, it is important to detect and analyze concept drift and minimize the performance degradation caused by it [257, 281]. In the field of machine learning, three main research topics, drift detection, drift understanding, and drift adaptation, are developed to analyze concept drift in streaming data. Machine learning researchers propose many automatic algorithms to detect and adapt concept drift. Although these algorithms can improve the adaptability of learning models in an uncertain environment, they only provide a numerical value to measure the drift degree at each time. This makes it hard to understand why and where drift occurs. If the adaption algorithms fail to improve the model performance, the black-box behavior of the adaption models makes it difficult to diagnose the root cause of performance degradation. As a result, model developers need to have tools that intuitively illustrate how data distributions have changed over time, which samples cause drift, and how the training samples and models can be adjusted for overcoming such drift. This requirement naturally leads to a visual analytics paradigm where the expert interacts and collaborates the concept drift detection and adaption algorithm by putting human in the loop. The major challenges here are 1) how to visually represent the evolution patterns of streaming data over time and effectively compare data distributions between/among time points; 2) tightly integrate such streaming data visualization with drift detection and adaption algorithms to form an interactive and progressive analysis environment with human in the loop.

7 Conclusions

In this paper, we comprehensively review recent progress and developments of visual analytics techniques for machine learning. These techniques are classified into three groups by the corresponding analysis stage: techniques before model building, techniques in model building, and techniques after model building. Each category is detailed by typical analysis tasks, and each task is featured by a set of representative works. By comprehensively analyzing existing visual analytics research for machine learning, we also suggest six directions in future machine-learning-related visual analytics research, including improving data quality for weakly supervised learning and explainable feature engineering before model building, online training diagnosis and intelligent model refinement in model building, as well as multi-modal data understanding and concept drift analysis after model building. We hope this survey can provide an overview of the visual analytics research for machine learning, facilitate the understanding of state-of-the-art knowledge in this area, and shed light on future research.

References

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean,

M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[2] Y. Ahn and Y.-R. Lin. FairSight: Visual analytics for fairness in decision making. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1086–1095, 2020.

[3] S. Alemzadeh, U. Niemann, T. Ittermann, H. Völzke, D. Schneider, M. Spiliopoulou, K. Bühler, and B. Preim. Visual analysis of missing values in longitudinal cohort study data. *Computer Graphics Forum*, 39(1):63–75, 2020.

[4] E. Alexander, J. Kohlmann, R. Valenza, M. Witmore, and M. Gleicher. Serendip: Topic model-driven visual exploration of text corpora. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 173–182, 2014.

[5] J. Alsakran, Y. Chen, D. Luo, Y. Zhao, J. Yang, W. Dou, and S. Liu. Real-time visualization of streaming text with a force-based dynamic system. *IEEE Computer Graphics and Applications*, 32(1):34–45, 2012.

[6] J. Alsakran, Y. Chen, Y. Zhao, J. Yang, and D. Luo. STREAMIT: Dynamic visualization and interactive exploration of text streams. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 131–138, 2011.

[7] B. Alsallakh, A. Hanbury, H. Hauser, S. Miksch, and A. Rauber. Visual methods for analyzing probabilistic classification data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1703–1712, 2014.

[8] G. Andrienko, N. Andrienko, G. Anzer, P. Bauer, G. Budziak, G. Fuchs, D. Hecker, H. Weber, and S. Wrobel. Constructing spaces and times for tactical analysis in football. *IEEE Transactions on Visualization and Computer Graphics*, 2019. to be published, doi: 10.1109/TVCG.2019.2952129.

[9] G. Andrienko, N. Andrienko, S. Bremm, T. Schreck, T. Von Landesberger, P. Bak, and D. Keim. Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns. *Computer Graphics Forum*, 29(3):913–922, 2010.

[10] G. Andrienko, N. Andrienko, C. Hurter, S. Rinzivillo, and S. Wrobel. Scalable analysis of movement data for extracting and exploring significant places. *IEEE Transactions on Visualization and Computer Graphics*, 19(7):1078–1094, 2013.

[11] C. Arbesser, F. Spechtenhauser, T. Mühlbacher, and H. Piringer. Visplause: Visual data quality assessment of many time series using plausibility checks. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):641–650, 2016.

[12] B. O. Ayinde and J. M. Zurada. Building efficient convnets using redundant feature pruning. In *Proceedings of the International Conference on Learning Representations*, 2018.

[13] T. Baltrušaitis, C. Ahuja, and L.-P. Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, 2018.

[14] A. Bäuerle, H. Neumann, and T. Ropinski. Classifier-guided visual correction of noisy labels for image classification tasks. *Computer Graphics Forum*, 2020. to be published, doi: 10.1111/cgf.13973.

[15] M. Berger, K. McDonough, and L. M. Seversky. cite2vec: Citation-driven document exploration via word embeddings. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):691–700, 2017.

[16] J. Bernard, M. Hutter, H. Reinemuth, H. Pfeifer, C. Bors, and J. Kohlhammer. Visual-interactive preprocessing of multivariate time series data. *Computer Graphics Forum*, 38(3):401–412, 2019.

[17] J. Bernard, M. Hutter, M. Zeppelzauer, D. Fellner, and M. Sedlmair. Comparing visual-interactive labeling with active learning: An experimental study. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):298–308, 2017.

[18] J. Bernard, M. Zeppelzauer, M. Lehmann, M. Müller, and M. Sedlmair. Towards user-centered active learning algorithms. *Computer Graphics Forum*, 37(3):121–132, 2018.

[19] A. Bilal, A. Jourabloo, M. Ye, X. Liu, and L. Ren. Do convolutional neural networks learn class hierarchy? *IEEE Transactions on Visualization and Computer Graphics*, 24(1):152–162, 2018.

[20] T. Blascheck, F. Beck, S. Baltes, T. Ertl, and D. Weiskopf. Visual analysis and coding of data-rich user behavior. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 141–150, 2016.

[21] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022, 2003.

[22] M. Blumenschein, M. Behrisch, S. Schmid, S. Butscher, D. R. Wahl, K. Villinger, B. Renner, H. Reiterer, and D. A. Keim. SMARTexplore: Simplifying high-dimensional data analysis through a table-based visual analytics approach. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 36–47, 2018.

[23] M. Bogl, W. Aigner, P. Filzmoser, T. Lammarsch, S. Miksch, and A. Rind. Visual analytics for model selection in time series analysis. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2237–2246, 2013.

[24] M. Bögl, P. Filzmoser, T. Gschwandtner, T. Lammarsch, R. A. Leite, S. Miksch, and A. Rind. Cycle plot revisited: Multivariate outlier detection using a distance-based abstraction. *Computer Graphics Forum*, 36(3):227–238, 2017.

[25] C. Bors, T. Gschwandtner, and S. Miksch. Capturing and visualizing provenance from data wrangling. *IEEE Computer Graphics and Applications*, 39(6):61–75, 2019.

[26] H. Bosch, D. Thom, F. Heimerl, E. Püttmann, S. Koch, R. Krüger, M. Wörner, and T. Ertl. Scatterblogs2: Real-time monitoring of microblog messages through user-guided filtering. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2022–2031, 2013.

[27] L. Bradel, C. North, and L. House. Multi-model semantic interaction for text analytics. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 163–172, 2014.

[28] B. Broeksema, T. Baudel, A. Telea, and P. Crisafulli. Decision exploration lab: A visual analytics solution for decision management. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):1972–1981, 2013.

[29] B. Broeksema, A. C. Telea, and T. Baudel. Visual analysis of multi-dimensional categorical data sets. *Computer Graphics Forum*, 32(8):158–169, 2013.

[30] M. Brooks, S. Amershi, B. Lee, S. M. Drucker, A. Kapoor, and P. Simard. FeatureInsight: Visual support for error-driven feature ideation in text classification. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 105–112, 2015.

[31] J. Buchmüller, H. Janetzko, G. Andrienko, N. Andrienko, G. Fuchs, and D. A. Keim. Visual analytics for exploring local impact of air traffic. *Computer Graphics Forum*, 34(3):181–190, 2015.

[32] A. A. Cabrera, W. Epperson, F. Hohman, M. Kahng, J. Morgenstern, and D. H. Chau. FAIRVIS: Visual analytics for discovering intersectional bias in machine learning. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 46–56, 2019.

[33] K. Cao, M. Liu, H. Su, J. Wu, J. Zhu, and S. Liu. Analyzing the noise robustness of deep neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2020.2969185.

[34] N. Cao, C. Lin, Q. Zhu, Y.-R. Lin, X. Teng, and X. Wen. Voila: Visual anomaly detection and monitoring with streaming spatiotemporal data. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):23–33, 2018.

[35] N. Cao, Y.-R. Lin, X. Sun, D. Lazer, S. Liu, and H. Qu. Whisper: Tracing the spatiotemporal process of information diffusion in real time. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2649–2658, 2012.

[36] N. Cao, C. Shi, S. Lin, J. Lu, Y.-R. Lin, and C.-Y. Lin. TargetVue: Visual analysis of anomalous user behaviors in online communication systems. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):280–289, 2016.

[37] N. Cao, J. Sun, Y.-R. Lin, D. Gotz, S. Liu, and H. Qu. Facetatlas: Multifaceted visualization for rich text corpora. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1172–1181, 2010.

[38] D. Cashman, G. Patterson, A. Mosca, N. Watts, S. Robinson, and R. Chang. RNNbow: Visualizing learning via backpropagation gradients in RNNs. *IEEE Computer Graphics and Applications*, 38(6):39–50, 2018.

[39] D. Cashman, A. Perer, R. Chang, and H. Strobelt. Ablate, variate, and contemplate: Visual analytics for discovering neural architectures. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):863–873, 2020.

[40] M. Cavallo and Ç. Demiralp. Track xplorer: A system for visual analysis of sensor-based motor activity predictions. *Computer Graphics Forum*, 37(3):339–349, 2018.

[41] M. Cavallo and C. Demiralp. Clustrophile 2: Guided visual clustering analysis. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):267–276, 2019.

[42] J. Chae, D. Thom, H. Bosch, Y. Jang, R. Maciejewski, D. S. Ebert, and T. Ertl. Spatiotemporal social media analytics for abnormal event detection and examination using seasonal-trend decomposition. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 143–152, 2012.

[43] S. Chandrasegaran, S. K. Badam, L. Kisselburgh, K. Ramani, and N. Elmqvist. Integrating visual analytics support for grounded theory practice in qualitative text analysis. *Computer Graphics Forum*, 36(3):201–212, 2017.

[44] G. Chandrashekar and F. Sahin. A survey on feature selection methods. *Computers & Electrical Engineering*, 40(1):16–28, 2014.

[45] C. Chen, J. Yuan, Y. Lu, Y. Liu, H. Su, S. Yuan, and S. Liu. OoDAnalyzer: Interactive analysis of out-of-distribution samples. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2020.2973258.

[46] Q. Chen, X. Yue, X. Plantaz, Y. Chen, C. Shi, T.-C. Pong, and H. Qu. ViSeq: Visual analytics of learning sequence in massive open online courses. *IEEE Transactions on Visualization and Computer Graphics*, 26(3):1622–1636, 2020.

[47] S. Chen, N. Andrienko, G. Andrienko, L. Adilova, J. Barlet, J. Kindermann, P. H. Nguyen, O. Thonnard, and C. Turkay. LDA ensembles for interactive exploration and categorization of

behaviors. *IEEE Transactions on Visualization and Computer Graphics*, 2019. to be published, doi: 10.1109/TVCG.2019.2904069.

[48] S. Chen, S. Chen, L. Lin, X. Yuan, J. Liang, and X. Zhang. E-map: A visual analytics approach for exploring significant event evolutions in social media. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 36–47, 2017.

[49] S. Chen, S. Chen, Z. Wang, J. Liang, X. Yuan, N. Cao, and Y. Wu. D-Map: Visual analysis of ego-centric information diffusion patterns in social media. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 41–50, 2016.

[50] S. Chen, X. Yuan, Z. Wang, C. Guo, J. Liang, Z. Wang, X. Zhang, and J. Zhang. Interactive visual discovering of movement patterns from sparsely sampled geo-tagged social media data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):270–279, 2016.

[51] Y. Chen, Q. Chen, M. Zhao, S. Boyer, K. Veeramachaneni, and H. Qu. 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*, pages 111–120, 2016.

[52] Y. Chen, P. Xu, and L. Ren. Sequence synopsis: Optimize visual summary of temporal event data. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):45–55, 2018.

[53] J. Choo, C. Lee, C. K. Reddy, and H. Park. UTOPIAN: User-driven topic modeling based on interactive nonnegative matrix factorization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):1992–2001, 2013.

[54] J. Choo and S. Liu. Visual analytics for explainable deep learning. *IEEE Computer Graphics and Applications*, 38(4):84–92, 2018.

[55] D. Chu, D. A. Sheets, Y. Zhao, Y. Wu, J. Yang, M. Zheng, and G. Chen. Visualizing hidden themes of taxi movement with semantic transformation. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 137–144, 2014.

[56] D. Collaris and J. J. van Wijk. ExplainExplore: Visual exploration of machine learning explanations. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 26–35, 2020.

[57] M. Correll, M. Witmore, and M. Gleicher. Exploring collections of tagged text for literary scholarship. *Computer Graphics Forum*, 30(3):731–740, 2011.

[58] W. Cui, S. Liu, L. Tan, C. Shi, Y. Song, Z. Gao, H. Qu, and X. Tong. TextFlow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2412–2421, 2011.

[59] W. Cui, S. Liu, Z. Wu, and H. Wei. How hierarchical topics evolve in large text corpora. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2281–2290, 2014.

[60] S. Das, D. Cashman, R. Chang, and A. Endert. BEAMES: Interactive multimodel steering, selection, and inspection for regression tasks. *IEEE Computer Graphics and Applications*, 39(5):20–32, 2019.

[61] K. Dextras-Romagnino and T. Munzner. Segmentifier: Interactive refinement of clickstream data. *Computer Graphics Forum*, 38(3):623–634, 2019.

[62] G. Di Lorenzo, M. Sbodio, F. Calabrese, M. Berlingerio, F. Pinelli, and R. Nair. AllAboard: Visual exploration of cellphone mobility data to optimise public transport. *IEEE Transactions on Visualization and Computer Graphics*, 22(2):1036–1050, 2016.

[63] A. Diehl, L. Pelorosso, C. Delrieux, K. Matković, J. Ruiz, M. Gröller, and S. Bruckner. Albero: A visual analytics approach for probabilistic weather forecasting. *Computer Graphics Forum*, 36(7):135–144, 2017.

[64] D. Dingen, M. vant Veer, P. Houthuizen, E. H. J. Mestrom, E. H. Korsten, A. R. Bouwman, and J. van Wijk. RegressionExplorer: Interactive exploration of logistic regression models with subgroup analysis. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):246–255, 2019.

[65] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. DeCAF: A deep convolutional activation feature for generic visual recognition. In *Proceedings of the International Conference on Machine Learning*, pages 647–655, 2014.

[66] W. Dou, I. Cho, O. ElTayeby, J. Choo, X. Wang, and W. Ribarsky. DemographicVis: Analyzing demographic information based on user generated content. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 57–64, 2015.

[67] W. Dou, X. Wang, R. Chang, and W. Ribarsky. ParallelTopics: A probabilistic approach to exploring document collections. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 231–240, 2011.

[68] W. Dou, X. Wang, D. Skau, W. Ribarsky, and M. X. Zhou. Leadline: Interactive visual analysis of text data through event identification and exploration. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 93–102, 2012.

[69] W. Dou, L. Yu, X. Wang, Z. Ma, and W. Ribarsky. HierarchicalTopics: Visually exploring large text collections using topic hierarchies. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2002–2011, 2013.

[70] F. Du, C. Plaisant, N. Spring, and B. Shneiderman. EventAction: Visual analytics for temporal event

sequence recommendation. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 61–70, 2016.

[71] C. Eichner, H. Schumann, and C. Tominski. Making parameter dependencies of time-series segmentation visually understandable. *Computer Graphics Forum*, 39(1):607–622, 2019.

[72] M. El-Assady, V. Gold, C. Acevedo, C. Collins, and D. Keim. ConToVi: Multi-party conversation exploration using topic-space views. *Computer Graphics Forum*, 35(3):431–440, 2016.

[73] M. El-Assady, R. Kehlbeck, C. Collins, D. Keim, and O. Deussen. Semantic concept spaces: Guided topic model refinement using word-embedding projections. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1001–1011, 2020.

[74] M. El-Assady, R. Sevastjanova, B. Gipp, D. Keim, and C. Collins. NEREx: Named-entity relationship exploration in multi-party conversations. *Computer Graphics Forum*, 36(3):213–225, 2017.

[75] M. El-Assady, R. Sevastjanova, D. Keim, and C. Collins. ThreadReconstructor: Modeling reply-chains to untangle conversational text through visual analytics. *Computer Graphics Forum*, 37(3):351–365, 2018.

[76] M. El-Assady, R. Sevastjanova, F. Sperrle, D. Keim, and C. Collins. Progressive learning of topic modeling parameters: A visual analytics framework. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):382–391, 2018.

[77] M. El-Assady, F. Sperrle, O. Deussen, D. Keim, and C. Collins. Visual analytics for topic model optimization based on user-steerable speculative execution. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):374–384, 2019.

[78] M. Fan, K. Wu, J. Zhao, Y. Li, W. Wei, and K. N. Truong. VisTA: Integrating machine intelligence with visualization to support the investigation of think-aloud sessions. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):343–352, 2020.

[79] N. Ferreira, L. Lins, D. Fink, S. Kelling, C. Wood, J. Freire, and C. Silva. BirdVis: Visualizing and understanding bird populations. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2374–2383, 2011.

[80] N. Ferreira, J. Poco, H. T. Vo, J. Freire, and C. T. Silva. Visual exploration of big spatio-temporal urban data: A study of new york city taxi trips. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2149–2158, 2013.

[81] V. Filipov, A. Arleo, P. Federico, and S. Miksch. CV3: visual exploration, assessment, and comparison of cvs. *Computer Graphics Forum*, 38(3):107–118, 2019.

[82] J. Foulds and E. Frank. A review of multi-instance learning assumptions. *The Knowledge Engineering Review*, 25(1):1–25, 2010.

[83] D. Fried and S. G. Kobourov. Maps of computer science. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 113–120, 2014.

[84] B. Fröhler, T. Möller, and C. Heinzl. GEMSe: Visualization-guided exploration of multi-channel segmentation algorithms. *Computer Graphics Forum*, 35(3):191–200, 2016.

[85] J. Fulda, M. Brehmer, and T. Munzner. TimeLineCurator: Interactive authoring of visual timelines from unstructured text. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):300–309, 2016.

[86] M. Gleicher, A. Barve, X. Yu, and F. Heimerl. Boxer: Interactive Comparison of Classifier Results. *Computer Graphics Forum*, 2020. to be published, doi: 10.1111/cgf.13972.

[87] M. Glueck, M. P. Naeini, F. Doshi-Velez, F. Chevalier, A. Khan, D. Wigdor, and M. Brudno. PhenoLines: phenotype comparison visualizations for disease subtyping via topic models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):371–381, 2018.

[88] B. Gobbo, D. Balsamo, M. Mauri, P. Bajardi, A. Panisson, and P. Ciuccarelli. Topic Tomographies (TopTom): a visual approach to distill information from media streams. *Computer Graphics Forum*, 38(3):609–621, 2019.

[89] C. Görg, Z. Liu, J. Kihm, J. Choo, H. Park, and J. Stasko. Combining computational analyses and interactive visualization for document exploration and sensemaking in jigsaw. *IEEE Transactions on Visualization and Computer Graphics*, 19(10):1646–1663, 2013.

[90] D. Gotz and H. Stavropoulos. DecisionFlow: Visual analytics for high-dimensional temporal event sequence data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1783–1792, 2014.

[91] T. Gschwandtner and O. Erhart. Know your enemy: Identifying quality problems of time series data. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 205–214, 2018.

[92] H. Guo and D. H. Laidlaw. Topic-based exploration and embedded visualizations for research idea generation. *IEEE Transactions on Visualization and Computer Graphics*, 26(3):1592–1607, 2020.

[93] S. Guo, Z. Jin, D. Gotz, F. Du, H. Zha, and N. Cao. Visual progression analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):417–426, 2019.

[94] S. Guo, K. Xu, R. Zhao, D. Gotz, H. Zha, and N. Cao. EventThread: Visual summarization and stage analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):56–65, 2018.

[95] I. Gutenko, K. Dmitriev, A. E. Kaufman, and M. A. Barish. AnaFe: Visual analytics of image-derived temporal features—focusing on the spleen. *IEEE Transactions on Visualization and Computer*

Graphics, 23(1):171–180, 2017.

[96] G. Halter, R. Ballester-Ripoll, B. Flueckiger, and R. Pajarola. VIAN: A visual annotation tool for film analysis. *Computer Graphics Forum*, 38(3):119–129, 2019.

[97] S. Havre, E. Hetzler, P. Whitney, and L. Nowell. ThemeRiver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):9–20, 2002.

[98] W. He, T.-Y. Lee, J. van Baar, K. Wittenburg, and H.-W. Shen. DynamicsExplorer: Visual analytics for robot control tasks involving dynamics and LSTM-based control policies. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 36–45, 2020.

[99] F. Heimerl, Q. Han, S. Koch, and T. Ertl. CiteRivers: Visual analytics of citation patterns. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):190–199, 2016.

[100] F. Heimerl, M. John, Q. Han, S. Koch, and T. Ertl. DocuCompass: Effective exploration of document landscapes. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 11–20, 2016.

[101] F. Heimerl, S. Koch, H. Bosch, and T. Ertl. Visual classifier training for text document retrieval. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2839–2848, 2012.

[102] B. Höferlin, R. Netzel, M. Höferlin, D. Weiskopf, and G. Heidemann. Inter-active learning of ad-hoc classifiers for video visual analytics. In *Proceedings of the Conference on Visual Analytics Science and Technology*, pages 23–32, 2012.

[103] F. Hohman, M. Kahng, R. Pienta, and D. H. Chau. Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE Transactions on Visualization and Computer Graphics*, 25(8):2674–2693, 2019.

[104] F. Hohman, H. Park, C. Robinson, and D. H. P. Chau. Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1096–1106, 2020.

[105] F. Hong, C. Lai, H. Guo, E. Shen, X. Yuan, and S. Li. FLDA: Latent dirichlet allocation based unsteady flow analysis. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2545–2554, 2014.

[106] E. Hoque and G. Carenini. Convis: A visual text analytic system for exploring blog conversations. *Computer Graphics Forum*, 33(3):221–230, 2014.

[107] M. Hu, K. Wongsuphasawat, and J. Stasko. Visualizing social media content with sententree. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):621–630, 2017.

[108] N. Q. V. Hung, D. C. Thang, M. Weidlich, and K. Aberer. Minimizing efforts in validating crowd answers. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 999–1014, 2015.

[109] S. Ingram, T. Munzner, V. Irvine, M. Tory, S. Bergner, and T. Möller. Dimstiller: Workflows for dimensional analysis and reduction. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 3–10, 2010.

[110] M. Itoh, M. Toyoda, C. Z. Zhu, S. Satoh, and M. Kitsuregawa. Image flows visualization for inter-media comparison. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 129–136, 2014.

[111] M. Itoh, N. Yoshinaga, M. Toyoda, and M. Kitsuregawa. Analysis and visualization of temporal changes in bloggers' activities and interests. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 57–64, 2012.

[112] H. Jänicke, R. Borgo, J. S. Mason, and M. Chen. SoundRiver: Semantically-rich sound illustration. *Computer Graphics Forum*, 29(2):357–366, 2010.

[113] S. Jänicke and D. J. Wrisley. Interactive visual alignment of medieval text versions. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 127–138, 2017.

[114] M. Jankowska, V. Kešelj, and E. Milios. 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*, pages 103–112, 2012.

[115] T. Jaunet, R. Vuillemot, and C. Wolf. DRLViz: Understanding decisions and memory in deep reinforcement learning. *Computer Graphics Forum*, 2020. to be published, doi: 10.1111/cgf.13962.

[116] C. S. Jean, C. Ware, and R. Gamble. Dynamic change arcs to explore model forecasts. *Computer Graphics Forum*, 35(3):311–320, 2016.

[117] X. Ji, H.-W. Shen, A. Ritter, R. Machiraju, and P.-Y. Yen. Visual exploration of neural document embedding in information retrieval: semantics and feature selection. *IEEE Transactions on Visualization and Computer Graphics*, 25(6):2181–2192, 2019.

[118] A. S. Júnior, C. Renso, and S. Matwin. ANALYTiC: An active learning system for trajectory classification. *IEEE Computer Graphics and Applications*, 37(5):28–39, 2017.

[119] M. Kahng, P. Y. Andrews, A. Kalro, and D. H. Chau. ActiVis: Visual exploration of industry-scale deep neural network models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):88–97, 2018.

[120] M. Kahng, N. Thorat, D. H. P. Chau, F. B. Viegas, and M. Wattenberg. GAN lab: Understanding complex deep generative models using interactive visual experimentation. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):310–320, 2019.

[121] T. Kakar, X. Qin, E. A. Rundensteiner, L. Harrison,

S. K. Sahoo, and S. De. DIVA: Exploration and validation of hypothesized drug-drug interactions. *Computer Graphics Forum*, 38(3):95–106, 2019.

[122] R. Kamaleswaran, C. Collins, A. James, and C. McGregor. PhysioEx: visual analysis of physiological event streams. *Computer Graphics Forum*, 35(3):331–340, 2016.

[123] S. Kandel, R. Parikh, A. Paepcke, J. M. Hellerstein, and J. Heer. Profiler: Integrated statistical analysis and visualization for data quality assessment. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, pages 547–554, 2012.

[124] A. Karduni, I. Cho, G. Wessel, W. Ribarsky, E. Sauda, and W. Dou. Urban space explorer: a visual analytics system for urban planning. *IEEE Computer Graphics and Applications*, 37(5):50–60, 2017.

[125] M. Khayat, M. Karimzadeh, J. Zhao, and D. S. Ebert. VASSL: A visual analytics toolkit for social spambot labeling. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):874–883, 2019.

[126] H. Kim, D. Choi, B. Drake, A. Endert, and H. Park. TopicSifter: Interactive search space reduction through targeted topic modeling. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 35–45, 2019.

[127] H. Kim, B. Drake, A. Endert, and H. Park. ArchiText: Interactive hierarchical topic modeling. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2020.2981456.

[128] M. Kim, K. Kang, D. Park, J. Choo, and N. Elmqvist. TopicLens: Efficient multi-level visual topic exploration of large-scale document collections. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):151–160, 2017.

[129] A. Kochtchi, T. v. Landesberger, and C. Biemann. Networks of names: Visual exploration and semi-automatic tagging of social networks from newspaper articles. *Computer Graphics Forum*, 33(3):211–220, 2014.

[130] R. Kosara, F. Bendix, and H. Hauser. Parallel sets: Interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):558–568, 2006.

[131] J. Krause, A. Dasgupta, J. Swartz, Y. Aphinyanaphongs, and E. Bertini. A workflow for visual diagnostics of binary classifiers using instance-level explanations. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 162–172, 2017.

[132] J. Krause, A. Perer, and E. Bertini. INFUSE: Interactive feature selection for predictive modeling of high dimensional data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1614–1623, 2014.

[133] R. Krueger, Q. Han, N. Ivanov, S. Mahtal, D. Thom, H. Pfister, and T. Ertl. Bird's-eye-large-scale visual analytics of city dynamics using social location data. *Computer Graphics Forum*, 38(3):595–607, 2019.

[134] R. Krueger, D. Thom, and T. Ertl. Visual analysis of movement behavior using web data for context enrichment. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 193–200, 2014.

[135] R. Krüger, D. Thom, and T. Ertl. Semantic enrichment of movement behavior with foursquare–a visual analytics approach. *IEEE Transactions on Visualization and Computer Graphics*, 21(8):903–915, 2015.

[136] K. Kurzhals, M. Hlawatsch, C. Seeger, and D. Weiskopf. Visual analytics for mobile eye tracking. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):301–310, 2016.

[137] B. C. Kwon, V. Anand, K. A. Severson, S. Ghosh, Z. Sun, B. I. Frohnert, M. Lundgren, and K. Ng. DPVis: Visual analytics with hidden markov models for disease progression pathways. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2020.2985689.

[138] B. C. Kwon, M.-J. Choi, J. T. Kim, E. Choi, Y. B. Kim, S. Kwon, J. Sun, and J. Choo. RetainVis: Visual analytics with interpretable and interactive recurrent neural networks on electronic medical records. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):299–309, 2019.

[139] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 6402–6413, 2017.

[140] C. Lee, Y. Kim, S. M. Jin, D. Kim, R. Maciejewski, D. Ebert, and S. Ko. A visual analytics system for exploring, monitoring, and forecasting road traffic congestion. *IEEE Transactions on Visualization and Computer Graphics*, 2019. to be published, doi: 10.1109/TVCG.2019.2922597.

[141] H. Lee, J. Kihm, J. Choo, J. Stasko, and H. Park. iVisClustering: An interactive visual document clustering via topic modeling. *Computer Graphics Forum*, 31(3):1155–1164, 2012.

[142] K. Lee, H. Lee, K. Lee, and J. Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. In *Proceedings of the International Conference on Learning Representations*, 2018.

[143] R. A. Leite, T. Gschwandtner, S. Miksch, S. Kriglstein, M. Pohl, E. Gstrein, and J. Kuntner. EVA: Visual analytics to identify fraudulent events. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):330–339, 2018.

[144] F. Lekschas, B. Peterson, D. Haehn, E. Ma, N. Gehlenborg, and H. Pfister. PEAX: Interactive visual pattern search in sequential data using unsupervised deep representation learning. *Computer Graphics Forum*, 2020. to be published, doi:

10.1111/cgf.13971.
[145] J. Li, S. Chen, W. Chen, G. Andrienko, and N. Andrienko. Semantics-space-time cube. a conceptual framework for systematic analysis of texts in space and time. *IEEE Transactions on Visualization and Computer Graphics*, 26(4):1789–1806, 2020.
[146] M. Li, F. Choudhury, Z. Bao, H. Samet, and T. Sellis. ConcaveCubes: Supporting cluster-based geographical visualization in large data scale. *Computer Graphics Forum*, 37(3):217–228, 2018.
[147] Q. Li, Z. Wu, L. Yi, K. S. N., H. Qu, and X. Ma. WeSeer: Visual analysis for better information cascade prediction of WeChat articles. *IEEE Transactions on Visualization and Computer Graphics*, 26(2):1399–1412, 2020.
[148] Y.-F. Li, L.-Z. Guo, and Z.-H. Zhou. Towards safe weakly supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. to be published, doi: 10.1109/TPAMI.2019.2922396.
[149] Y.-F. Li, S.-B. Wang, and Z.-H. Zhou. Graph quality judgement: A large margin expedition. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1725–1731, 2016.
[150] Z. Li, C. Zhang, S. Jia, and J. Zhang. Galex: Exploring the evolution and intersection of disciplines. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1182–1192, 2020.
[151] H. Liu, S. Jin, Y. Yan, Y. Tao, and H. Lin. Visual analytics of taxi trajectory data via topical sub-trajectories. *Visual Informatics*, 3(3):140–149, 2019.
[152] M. Liu, L. Jiang, J. Liu, X. Wang, J. Zhu, and S. Liu. Improving learning-from-crowds through expert validation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 2329–2336, 2017.
[153] M. Liu, S. Liu, X. Zhu, Q. Liao, F. Wei, and S. Pan. An uncertainty-aware approach for exploratory microblog retrieval. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):250–259, 2016.
[154] M. Liu, J. Shi, K. Cao, J. Zhu, and S. Liu. Analyzing the training processes of deep generative models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):77–87, 2018.
[155] M. Liu, J. Shi, Z. Li, C. Li, J. Zhu, and S. Liu. Towards better analysis of deep convolutional neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):91–100, 2017.
[156] S. Liu, C. Chen, Y. Lu, F. Ouyang, and B. Wang. An interactive method to improve crowdsourced annotations. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):235–245, 2019.
[157] S. Liu, W. Cui, Y. Wu, and M. Liu. A survey on information visualization: recent advances and challenges. *The Visual Computer*, 30(12):1373–1393, 2014.
[158] S. Liu, Z. Li, T. Li, V. Srikumar, V. Pascucci, and P.-T. Bremer. NLIZE: A perturbation-driven visual interrogation tool for analyzing and interpreting natural language inference models. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):651–660, 2019.
[159] S. Liu, B. Wang, J. J. Thiagarajan, P.-T. Bremer, and V. Pascucci. Visual exploration of high-dimensional data through subspace analysis and dynamic projections. *Computer Graphics Forum*, 34(3):271–280, 2015.
[160] S. Liu, X. Wang, J. Chen, J. Zhu, and B. Guo. TopicPanorama: A full picture of relevant topics. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 183–192, 2014.
[161] S. Liu, X. Wang, C. Collins, W. Dou, F. Ouyang, M. El-Assady, L. Jiang, and D. A. Keim. Bridging text visualization and mining: A task-driven survey. *IEEE Transactions on Visualization and Computer Graphics*, 25(7):2482–2504, 2019.
[162] S. Liu, X. Wang, M. Liu, and J. Zhu. Towards better analysis of machine learning models: A visual analytics perspective. *Visual Informatics*, 1(1):48–56, 2017.
[163] S. Liu, J. Xiao, J. Liu, X. Wang, J. Wu, and J. Zhu. Visual diagnosis of tree boosting methods. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):163–173, 2018.
[164] S. Liu, J. Yin, X. Wang, W. Cui, K. Cao, and J. Pei. Online visual analytics of text streams. *IEEE Transactions on Visualization and Computer Graphics*, 22(11):2451–2466, 2016.
[165] S. Liu, M. X. Zhou, S. Pan, Y. Song, W. Qian, W. Cai, and X. Lian. TIARA: Interactive, topic-based visual text summarization and analysis. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(2):1–28, 2012.
[166] X. Liu, A. Xu, L. Gou, H. Liu, R. Akkiraju, and H.-W. Shen. SocialBrands: Visual analysis of public perceptions of brands on social media. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 71–80, 2016.
[167] Z. Liu, B. Kerr, M. Dontcheva, J. Grover, M. Hoffman, and A. Wilson. CoreFlow: Extracting and visualizing branching patterns from event sequences. *Computer Graphics Forum*, 36(3):527–538, 2017.
[168] Z. Liu, Y. Wang, M. Dontcheva, M. Hoffman, S. Walker, and A. Wilson. Patterns and sequences: Interactive exploration of clickstreams to understand common visitor paths. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):321–330, 2017.
[169] T. Lowe, E.-C. Forster, G. Albuquerque, J.-P. Kreiss, and M. Magnor. Visual analytics for development and evaluation of order selection criteria for autoregressive processes. *IEEE Transactions on Visualization and

Computer Graphics, 22(1):151–159, 2016.
[170] J. Lu, D. Batra, D. Parikh, and S. Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 13–23, 2019.
[171] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang. Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12):2346–2363, 2018.
[172] Y. Lu, R. Garcia, B. Hansen, M. Gleicher, and R. Maciejewski. The state-of-the-art in predictive visual analytics. *Computer Graphics Forum*, 36(3):539–562, 2017.
[173] Y. Lu, M. Steptoe, S. Burke, H. Wang, J.-Y. Tsai, H. Davulcu, D. Montgomery, S. R. Corman, and R. Maciejewski. Exploring evolving media discourse through event cueing. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):220–229, 2016.
[174] Y. Lu, F. Wang, and R. Maciejewski. Business intelligence from social media: A study from the vast box office challenge. *IEEE Computer Graphics and Applications*, 34(5):58–69, 2014.
[175] Y. Lu, H. Wang, S. Landis, and R. Maciejewski. A visual analytics framework for identifying topic drivers in media events. *IEEE Transactions on Visualization and Computer Graphics*, 24(9):2501–2515, 2018.
[176] D. Luo, J. Yang, M. Krstajic, W. Ribarsky, and D. Keim. EventRiver: Visually exploring text collections with temporal references. *IEEE Transactions on Visualization and Computer Graphics*, 18(1):93–105, 2012.
[177] Y. Ma, T. Xie, J. Li, and R. Maciejewski. Explaining vulnerabilities to adversarial machine learning through visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1075–1085, 2020.
[178] Z. Ma, W. Dou, X. Wang, and S. Akella. Tag-latent dirichlet allocation: Understanding hashtags and their relationships. In *Proceedings of the IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technologies*, pages 260–267, 2013.
[179] R. Maciejewski, R. Hafen, S. Rudolph, S. G. Larew, M. A. Mitchell, W. S. Cleveland, and D. S. Ebert. Forecasting hotspots—a predictive analytics approach. *IEEE Transactions on Visualization and Computer Graphics*, 17(4):440–453, 2011.
[180] J. MacInnes, S. Santosa, and W. Wright. Visual classification: Expert knowledge guides machine learning. *IEEE Computer Graphics and Applications*, 30(1):8–14, 2010.
[181] A. Malik, R. Maciejewski, S. Towers, S. McCullough, and D. S. Ebert. Proactive spatiotemporal resource allocation and predictive visual analytics for community policing and law enforcement. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1863–1872, 2014.
[182] S. Marsland. *Machine learning: an algorithmic perspective*. CRC press, 2015.
[183] T. May, A. Bannach, J. Davey, T. Ruppert, and J. Kohlhammer. Guiding feature subset selection with an interactive visualization. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 111–120, 2011.
[184] M. Migut, J. van Gemert, and M. Worring. Interactive decision making using dissimilarity to visually represented prototypes. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 141–149, 2011.
[185] M. Migut and M. Worring. Visual exploration of classification models for risk assessment. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 11–18, 2010.
[186] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 3111–3119, 2013.
[187] Y. Ming, S. Cao, R. Zhang, Z. Li, Y. Chen, Y. Song, and H. Qu. Understanding hidden memories of recurrent neural networks. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 13–24, 2017.
[188] Y. Ming, H. Qu, and E. Bertini. RuleMatrix: Visualizing and understanding classifiers with rules. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):342–352, 2019.
[189] Y. Ming, P. Xu, F. Cheng, H. Qu, and L. Ren. ProtoSteer: Steering deep sequence model with prototypes. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):238–248, 2020.
[190] Y. Ming, P. Xu, H. Qu, and L. Ren. Interpretable and steerable sequence learning via prototypes. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 903–913, 2019.
[191] F. Miranda, H. Doraiswamy, M. Lage, K. Zhao, B. Goncalves, L. Wilson, M. Hsieh, and C. T. Silva. Urban pulse: Capturing the rhythm of cities. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):791–800, 2017.
[192] J. Moehrmann, S. Bernstein, T. Schlegel, G. Werner, and G. Heidemann. Improving the usability of hierarchical representations for interactively labeling large image data sets. In *Proceedings of the International Conference on Human-Computer Interaction*, pages 618–627, 2011.
[193] T. Muhlbacher, L. Linhardt, T. Moller, and H. Piringer. TreePOD: Sensitivity-aware selection of pareto-optimal decision trees. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):174–183,

2018.

[194] T. Mühlbacher and H. Piringer. A partition-based framework for building and validating regression models. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):1962–1971, 2013.

[195] S. Murugesan, S. Malik, F. Du, E. Koh, and T. M. Lai. Deepcompare: Visual and interactive comparison of deep learning model performance. *IEEE Computer Graphics and Applications*, 39:47–59, 2019.

[196] A. Ng. Machine learning and AI via brain simulations. `https://ai.stanford.edu/~ang/slides/DeepLearning-Mar2013.pptx`, 2013. Last accessed 2020-05-24.

[197] S. Nie, C. Healey, K. Padia, S. Leeman-Munk, J. Benson, D. Caira, S. Sethi, and R. Devarajan. Visualizing deep neural networks for text analytics. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 180–189, 2018.

[198] N. J. Nilsson. Introduction to machine learning: An early draft of a proposed textbook. `https://ai.stanford.edu/~nilsson/MLBOOK.pdf`, 1996. Last accessed 2020-05-24.

[199] D. Oelke, H. Strobelt, C. Rohrdantz, I. Gurevych, and O. Deussen. Comparative exploration of document collections: a visual analytics approach. *Computer Graphics Forum*, 33(3):201–210, 2014.

[200] E. Packer, P. Bak, M. Nikkila, V. Polishchuk, and H. J. Ship. Visual analytics for spatial clustering: Using a heuristic approach for guided exploration. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2179–2188, 2013.

[201] J. G. S. Paiva, W. R. Schwartz, H. Pedrini, and R. Minghim. An approach to supporting incremental visual data classification. *IEEE Transactions on Visualization and Computer Graphics*, 21(1):4–17, 2014.

[202] D. Park, S. Kim, J. Lee, J. Choo, N. Diakopoulos, and N. Elmqvist. ConceptVector: text visual analytics via interactive lexicon building using word embedding. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):361–370, 2018.

[203] J. H. Park, S. Nadeem, S. Boorboor, J. Marino, and A. E. Kaufman. CMed: Crowd analytics for medical imaging data. *IEEE Transactions on Visualization and Computer Graphics*, 2019. to be published, doi: 10.1109/TVCG.2019.2953026.

[204] J. H. Park, S. Nadeem, S. Mirhosseini, and A. Kaufman. C^2A: Crowd consensus analytics for virtual colonoscopy. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 21–30, 2016.

[205] F. V. Paulovich, F. M. Toledo, G. P. Telles, R. Minghim, and L. G. Nonato. Semantic wordification of document collections. *Computer Graphics Forum*, 31(3pt3):1145–1153, 2012.

[206] N. Pezzotti, T. Hollt, J. V. Gemert, B. P. Lelieveldt, E. Eisemann, and A. Vilanova. DeepEyes: Progressive visual analytics for designing deep neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):98–108, 2018.

[207] H. Piringer, W. Berger, and J. Krasser. HyperMoVal: Interactive visual validation of regression models for real-time simulation. *Computer Graphics Forum*, 29(3):983–992, 2010.

[208] O. Purwantiningsih, A. Sallaberry, S. Andary, A. Seilles, and J. Azé. Visual analysis of body movement in serious games for healthcare. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 229–233, 2016.

[209] P. E. Rauber, S. G. Fadel, A. X. Falcao, and A. C. Telea. Visualizing the hidden activity of artificial neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):101–110, 2017.

[210] D. Ren, S. Amershi, B. Lee, J. Suh, and J. D. Williams. Squares: Supporting interactive performance analysis for multiclass classifiers. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):61–70, 2017.

[211] P. Riehmann, D. Kiesel, M. Kohlhaas, and B. Froehlich. Visualizing a thinker's life. *IEEE Transactions on Visualization and Computer Graphics*, 25(4):1803–1816, 2019.

[212] M. Rohlig, M. Luboschik, F. Kruger, T. Kirste, H. Schumann, M. Bogl, B. Alsallakh, and S. Miksch. Supporting activity recognition by visual analytics. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 41–48, 2015.

[213] O. Rooij, J. van Wijk, and M. Worring. MediaTable: Interactive categorization of multimedia collections. *IEEE Computer Graphics and Applications*, 30(5):42–51, 2010.

[214] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

[215] D. Sacha, F. Al-Masoudi, M. Stein, T. Schreck, D. A. Keim, G. Andrienko, and H. Janetzko. Dynamic visual abstraction of soccer movement. *Computer Graphics Forum*, 36(3):305–315, 2017.

[216] D. Sacha, M. Kraus, J. Bernard, M. Behrisch, T. Schreck, Y. Asano, and D. A. Keim. SOMFlow: Guided exploratory cluster analysis with self-organizing maps and analytic provenance. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):120–130, 2018.

[217] D. Sacha, M. Kraus, D. A. Keim, and M. Chen. VIS4ML: An ontology for visual analytics assisted machine learning. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):385–395, 2019.

[218] A. Sarikaya, M. Correli, J. M. Dinis, D. H. O'Connor, and M. Gleicher. Visualizing co-occurrence of events

in populations of viral genome sequences. *Computer Graphics Forum*, 35(3):151–160, 2016.

[219] R. Scheepens, S. Michels, H. van de Wetering, and J. J. van Wijk. Rationale visualization for safety and security. *Computer Graphics Forum*, 34(3):191–200, 2015.

[220] T. Schultz and G. L. Kindlmann. Open-box spectral clustering: Applications to medical image analysis. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2100–2108, 2013.

[221] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.

[222] J. Seo and B. Shneiderman. A rank-by-feature framework for interactive exploration of multidimensional data. *Information Visualization*, 4(2):96–113, 2005.

[223] Q. Shen, Y. Wu, Y. Jiang, W. Zeng, A. K. H. LAU, A. Vianova, and H. Qu. Visual interpretation of recurrent neural network on multi-dimensional time-series forecast. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 61–70, 2020.

[224] Q. Shen, W. Zeng, Y. Ye, S. M. Arisona, S. Schubiger, R. Burkhard, and H. Qu. StreetVizor: Visual exploration of human-scale urban forms based on street views. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):1004–1013, 2018.

[225] C. Shi, Y. Wu, S. Liu, H. Zhou, and H. Qu. LoyalTracker: Visualizing loyalty dynamics in search engines. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1733–1742, 2014.

[226] Q. shi Zhang and S. chun Zhu. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering*, 19(1):27–39, 2018.

[227] L. S. Snyder, Y.-S. Lin, M. Karimzadeh, D. Goldwasser, and D. S. Ebert. Interactive learning for identifying relevant tweets to support real-time situational awareness. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):558–568, 2019.

[228] F. Sperrle, R. Sevastjanova, R. Kehlbeck, and M. El-Assady. VIANA: Visual interactive annotation of argumentation. In *Proceedings of the Conference on Visual Analytics Science and Technology*, pages 11–22, 2019.

[229] T. Spinner, U. Schlegel, H. Schafer, and M. El-Assady. explAIner: A visual analytics framework for interactive and explainable machine learning. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1064–1074, 2020.

[230] M. Steiger, J. Bernard, S. Mittelstädt, H. Lücke-Tieke, D. Keim, T. May, and J. Kohlhammer. Visual analysis of time-series similarities for anomaly detection in sensor networks. *Computer Graphics Forum*, 33(3):401–410, 2014.

[231] M. Stein, H. Janetzko, T. Breitkreutz, D. Seebacher, T. Schreck, M. Grossniklaus, I. D. Couzin, and D. A. Keim. Director's cut: Analysis and annotation of soccer matches. *IEEE Computer Graphics and Applications*, 36(5):50–60, 2016.

[232] L. Stopar, P. Skraba, M. Grobelnik, and D. Mladenic. StreamStory: exploring multivariate time series on multiple scales. *IEEE Transactions on Visualization and Computer Graphics*, 25(4):1788–1802, 2019.

[233] H. Strobelt, S. Gehrmann, M. Behrisch, A. Perer, H. Pfister, and A. M. Rush. Seq2seq-Vis: A visual debugging tool for sequence-to-sequence models. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):353–363, 2019.

[234] H. Strobelt, S. Gehrmann, H. Pfister, and A. M. Rush. LSTMVis: A tool for visual analysis of hidden state dynamics in recurrent neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):667–676, 2018.

[235] N. Sultanum, D. Singh, M. Brudno, and F. Chevalier. Doccurate: A curation-based approach for clinical text visualization. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):142–151, 2019.

[236] G. Sun, Y. Wu, S. Liu, T.-Q. Peng, J. J. Zhu, and R. Liang. EvoRiver: Visual analysis of topic coopetition on social media. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1753–1762, 2014.

[237] C.-Y. Sung, X.-Y. Huang, Y. Shen, F.-Y. Cherng, W.-C. Lin, and H.-C. Wang. Exploring online learners' interactive dynamics by visually analyzing their time-anchored comments. *Computer Graphics Forum*, 36(7):145–155, 2017.

[238] G. K. Tam, H. Fang, A. J. Aubrey, P. W. Grant, P. L. Rosin, D. Marshall, and M. Chen. Visualization of time-series data in parameter space for understanding facial dynamics. *Computer Graphics Forum*, 30(3):901–910, 2011.

[239] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Hierarchical dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006.

[240] D. Thom, H. Bosch, S. Koch, M. Wörner, and T. Ertl. Spatiotemporal anomaly detection through visual analysis of geolocated twitter messages. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 41–48, 2012.

[241] D. Thom, R. Krüger, and T. Ertl. Can Twitter save lives? A broad-scale study on visual social media analytics for public safety. *IEEE Transactions on Visualization and Computer Graphics*, 22(7):1816–1829, 2016.

[242] T. Tian and J. Zhu. Max-margin majority voting for learning from crowds. In *Proceedings of the Advances in Neural Information Processing Systems*,

pages 1621–1629, 2015.

[243] G. Tkachev, S. Frey, and T. Ertl. Local prediction models for spatiotemporal volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2019. to be published, doi: 10.1109/TVCG.2019.2961893.

[244] F.-Y. Tzeng and K.-L. Ma. Opening the black box - data driven visualization of neural networks. In *Proceedings of the IEEE Conference on Visualization*, pages 383–390, 2005.

[245] S. van den Elzen and J. J. van Wijk. BaobabView: Interactive construction and analysis of decision trees. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 151–160, 2011.

[246] C. Vehlow, F. Beck, P. Auwärter, and D. Weiskopf. Visualizing the evolution of communities in dynamic graphs. *Computer Graphics Forum*, 34(1):277–288, 2015.

[247] T. von Landesberger, D. Basgier, and M. Becker. Comparative local quality assessment of 3d medical image segmentations with focus on statistical shape model-based algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2537–2549, 2016.

[248] T. Von Landesberger, F. Brodkorb, P. Roskosch, N. Andrienko, G. Andrienko, and A. Kerren. MobilityGraphs: Visual analysis of mass mobility dynamics via spatio-temporal graphs and clustering. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):11–20, 2016.

[249] K. Vrotsou and A. Nordman. Exploratory visual sequence mining based on pattern-growth. *IEEE Transactions on Visualization and Computer Graphics*, 25(8):2597–2610, 2019.

[250] E. Wall, S. Das, R. Chawla, B. Kalidindi, E. T. Brown, and A. Endert. Podium: Ranking data using mixed-initiative visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):288–297, 2018.

[251] J. Wang, L. Gou, H.-W. Shen, and H. Yang. DQNViz: A visual analytics approach to understand deep q-networks. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):288–298, 2019.

[252] J. Wang, L. Gou, H. Yang, and H.-W. Shen. Ganviz: A visual analytics approach to understand the adversarial game. *IEEE Transactions on Visualization and Computer Graphics*, 24(6):1905–1917, 2018.

[253] J. Wang, L. Gou, W. Zhang, H. Yang, and H.-W. Shen. DeepVID: Deep visual interpretation and diagnosis for image classifiers via knowledge distillation. *IEEE Transactions on Visualization and Computer Graphics*, 25(6):2168–2180, 2019.

[254] J. Wang, W. Zhang, and H. Yang. SCANViz: Interpreting the symbol-concept association captured by deep neural networks through visual analytics. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 51–60, 2020.

[255] Q. Wang, J. Yuan, S. Chen, H. Su, H. Qu, and S. Liu. Visual genealogy of deep neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2019.2921323.

[256] X. Wang, W. Chen, J.-K. Chou, C. Bryan, H. Guan, W. Chen, R. Pan, and K.-L. Ma. GraphProtector: a visual interface for employing and assessing multiple privacy preserving graph algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):193–203, 2018.

[257] X. Wang, W. Chen, J. Xia, Z. Chen, D. Xu, X. Wu, M. Xu, and T. Schreck. Conceptexplorer: Visual analysis of concept drifts in multi-source time-series data, 2020. https://arxiv.org/abs/2007.15272.

[258] X. Wang, J.-K. Chou, W. Chen, H. Guan, W. Chen, T. Lao, and K.-L. Ma. A utility-aware visual approach for anonymizing multi-attribute tabular data. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):351–360, 2017.

[259] X. Wang, W. Dou, Z. Ma, J. Villalobos, Y. Chen, T. Kraft, and W. Ribarsky. I-SI: Scalable architecture for analyzing latent topical-level information from social media data. *Computer Graphics Forum*, 31(3pt4):1275–1284, 2012.

[260] X. Wang, S. Liu, Y. Chen, T.-Q. Peng, J. Su, J. Yang, and B. Guo. How ideas flow across multiple social groups. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 51–60, 2016.

[261] X. Wang, S. Liu, J. Liu, J. Chen, J. Zhu, and B. Guo. TopicPanorama: A full picture of relevant topics. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2508–2521, 2016.

[262] X. Wang, S. Liu, Y. Song, and B. Guo. Mining evolutionary multi-branch trees from text streams. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 722–730, 2013.

[263] Y. Wang, H. Haleem, C. Shi, Y. Wu, X. Zhao, S. Fu, and H. Qu. Towards easy comparison of local businesses using online reviews. *Computer Graphics Forum*, 37(3):63–74, 2018.

[264] F. Wei, S. Liu, Y. Song, S. Pan, M. X. Zhou, W. Qian, L. Shi, L. Tan, and Q. Zhang. TIARA: a visual exploratory text analytic system. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 153–162, 2010.

[265] J. Wei, Z. Shen, N. Sundaresan, and K.-L. Ma. Visual cluster exploration of web clickstream data. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 3–12, 2012.

[266] J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viegas, and J. Wilson. The

what-if tool: Interactive probing of machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):56–65, 2019.

[267] W. Willett, S. Ginosar, A. Steinitz, B. Hartmann, and M. Agrawala. Identifying redundancy and exposing provenance in crowdsourced data analysis. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2198–2206, 2013.

[268] K. Wongsuphasawat, D. Smilkov, J. Wexler, J. Wilson, D. Mane, D. Fritz, D. Krishnan, F. B. Viegas, and M. Wattenberg. Visualizing dataflow graphs of deep learning models in TensorFlow. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):1–12, 2018.

[269] A. Wu and H. Qu. Multimodal analysis of video collections: Visual exploration of presentation techniques in TED talks. *IEEE Transactions on Visualization and Computer Graphics*, 26(7):2429–2442, 2020.

[270] W. Wu, Y. Zheng, N. Cao, H. Zeng, B. Ni, H. Qu, and L. M. Ni. MobiSeg: Interactive region segmentation using heterogeneous mobility data. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 91–100, 2017.

[271] Y. Wu, Z. Chen, G. Sun, X. Xie, N. Cao, S. Liu, and W. Cui. StreamExplorer: A multi-stage system for visually exploring events in social streams. *IEEE Transactions on Visualization and Computer Graphics*, 24(10):2758–2772, 2018.

[272] Y. Wu, S. Liu, K. Yan, M. Liu, and F. Wu. OpinionFlow: Visual analysis of opinion diffusion on social media. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1763–1772, 2014.

[273] Y. Wu, N. Pitipornvivat, J. Zhao, S. Yang, G. Huang, and H. Qu. egoSlider: Visual analysis of egocentric network evolution. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):260–269, 2016.

[274] S. Xiang, X. Ye, J. Xia, J. Wu, Y. Chen, and S. Liu. Interactive correction of mislabeled training data. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 57–68, 2019.

[275] C. Xie, W. Chen, X. Huang, Y. Hu, S. Barlowe, and J. Yang. VAET: A visual analytics approach for e-transactions time-series. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1743–1752, 2014.

[276] X. Xie, X. Cai, J. Zhou, N. Cao, and Y. Wu. A semantic-based method for visualizing large image collections. *IEEE Transactions on Visualization and Computer Graphics*, 25(7):2362–2377, 2019.

[277] J. Xu, Y. Tao, H. Lin, R. Zhu, and Y. Yan. Exploring controversy via sentiment divergences of aspects in reviews. In *Proceedings of the IEEE Pacific Visualization Symposium*, pages 240–249, 2017.

[278] J. Xu, Y. Tao, Y. Yan, and H. Lin. Exploring evolution of dynamic networks via diachronic node embeddings. *IEEE Transactions on Visualization and Computer Graphics*, 26(7):2387–2402, 2020.

[279] P. Xu, H. Mei, L. Ren, and W. Chen. ViDX: Visual diagnostics of assembly line performance in smart factories. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):291–300, 2017.

[280] P. Xu, Y. Wu, E. Wei, T.-Q. Peng, S. Liu, J. J. Zhu, and H. Qu. Visual analysis of topic competition on social media. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2012–2021, 2013.

[281] W. Yang, Z. Li, M. Liu, Y. Lu, K. Cao, R. Maciejewski, and S. Liu. Diagnosing concept drift with visual analytics, 2020. https://arxiv.org/abs/2007.14372v2.

[282] W. Yang, X. Wang, J. Lu, W. Dou, and S. Liu. Interactive steering of hierarchical clustering. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2020.2995100.

[283] L. Yu, W. Wu, X. Li, G. Li, W. S. Ng, S.-K. Ng, Z. Huang, A. Arunan, and H. M. Watt. iVizTRANS: Interactive visual learning for home and work place detection from massive public transportation data. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 49–56, 2015.

[284] G. G. Zanabria, J. A. Silveira, J. Poco, A. Paiva, M. B. Nery, C. T. Silva, S. F. A. de Abreu, and L. G. Nonato. CrimAnalyzer: Understanding crime patterns in são paulo. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2019.2947515.

[285] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *Proceedings of the European Conference on Computer Vision*, pages 818–833, 2014.

[286] H. Zeng, X. Shu, Y. Wang, Y. Wang, L. Zhang, T.-C. Pong, and H. Qu. EmotionCues: Emotion-oriented visual summarization of classroom videos. *IEEE Transactions on Visualization and Computer Graphics*, 2020. to be published, doi: 10.1109/TVCG.2019.2963659.

[287] H. Zeng, X. Wang, A. Wu, Y. Wang, Q. Li, A. Endert, and H. Qu. EmoCo: Visual analysis of emotion coherence in presentation videos. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):927–937, 2020.

[288] W. Zeng, C.-W. Fu, S. Müller Arisona, A. Erath, and H. Qu. Visualizing waypoints-constrained origin-destination patterns for massive transportation data. *Computer Graphics Forum*, 35(8):95–107, 2016.

[289] C. Zhang, J. Yang, F. B. Zhan, X. Gong, J. D. Brender, P. H. Langlois, S. Barlowe, and Y. Zhao. 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*, pages 136–143, 2016.

[290] J. Zhang, B. Ahlbrand, A. Malik, J. Chae, Z. Min, S. Ko, and D. S. Ebert. A visual analytics framework for microblog data analysis at multiple scales of aggregation. *Computer Graphics Forum*, 35(3):441–450, 2016.

[291] J. Zhang, Y. Wang, P. Molino, L. Li, and D. S. Ebert. Manifold: A model-agnostic framework for interpretation and diagnosis of machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):364–373, 2019.

[292] J. Zhang, E. Yanli, J. Ma, Y. Zhao, B. Xu, L. Sun, J. Chen, and X. Yuan. Visual analysis of public utility service problems in a metropolis. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1843–1852, 2014.

[293] L. Zhang and H. Huang. Hierarchical narrative collage for digital photo album. *Computer Graphics Forum*, 31(7):2173–2181, 2012.

[294] J. Zhao, N. Cao, Z. Wen, Y. Song, Y.-R. Lin, and C. Collins. #FluxFlow: Visual analysis of anomalous information spreading on social media. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1773–1782, 2014.

[295] J. Zhao, F. Chevalier, C. Collins, and R. Balakrishnan. Facilitating discourse analysis with interactive visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2639–2648, 2012.

[296] K. Zhao, M. O. Ward, E. A. Rundensteiner, and H. N. Higgins. LoVis: Local pattern visualization for model refinement. *Computer Graphics Forum*, 33(3):331–340, 2014.

[297] X. Zhao, Y. Wu, D. L. Lee, and W. Cui. iForest: Interpreting random forests via visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):407–416, 2019.

[298] Y. Zhao, X. Luo, X. Lin, H. Wang, X. Kui, F. Zhou, J. Wang, Y. Chen, and W. Chen. Visual analytics for electromagnetic situation awareness in radio monitoring and management. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):590–600, 2020.

[299] Z. Zhou, L. Meng, C. Tang, Y. Zhao, Z. Guo, M. Hu, and W. Chen. Visual abstraction of large scale geospatial origin-destination movement data. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):43–53, 2019.

[300] Z. Zhou, Z. Ye, Y. Liu, F. Liu, Y. Tao, and W. Su. Visual analytics for spatial clusters of air-quality data. *IEEE Computer Graphics and Applications*, 37(5):98–105, 2017.

[301] Z.-H. Zhou. Multi-instance learning from supervised view. *Journal of Computer Science and Technology*, 21(5):800–809, 2006.

[302] Z.-H. Zhou. A brief introduction to weakly supervised learning. *National Science Review*, 5(1):44–53, 2018.

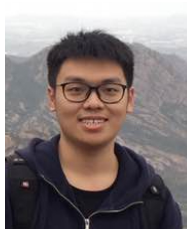

Jun Yuan is currently a Ph.D. student at Tsinghua University. His research interests are in explainable artificial intelligence. He received a B.S. degree from Tsinghua University.

Changjian Chen is now a Ph.D. student at Tsinghua University. His research interests are in interactive machine learning. He received a B.S. degree from University of Science and Technology of China.

Weikai Yang is a graduate student at Tsinghua University. His research interest is visual text analytics. He received a B.S. degree from Tsinghua University.

Mengchen Liu is a Senior Researcher at Microsoft. His research interest includes explainable AI and computer vision. He received a BS in Electronics Engineering and a Ph.D in Computer Science from Tsinghua University. He has served as PC members and reviewers in various conferences and journals.

Jiazhi Xia is an associate professor in School of Computer Science and Engineering at Central South University. He received his Ph.D. degree in Computer Science from The Nanyang Technological University, Singapore in 2011 and obtained his M.S. degree and B.S. degree in Computer Science and Technology from Zhejiang University in 2008 and 2005, respectively. His research interest includes data visualization, visual analytics, and computer graphics.

Shixia Liu is an associate professor at Tsinghua University. Her research interests include visual text analytics, visual social analytics, interactive machine learning, and text mining. She worked as a research staff member at IBM China Research Lab and a lead researcher at Microsoft Research Asia. She received a B.S. and M.S. from Harbin Institute of Technology, a Ph.D. from Tsinghua University. She is an associate editor-in-chief of IEEE Trans. Vis. Comput. Graph.