Computational Visual Media
https://doi.org/10.1007/s41095-020-0191-7 Vol. 7, No. 1, March 2021, 3–36

Review Article

A survey of visual analytics techniques for machine learning

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

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 to learn how to apply relevant techniques in visual analytics, we systematically review 259 papers published in the last ten years together with representative works before 2010. We build a taxonomy, which includes three first-level categories: techniques before model building, techniques during 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 promising potential future research opportunities 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 [1]. 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 has received ongoing and increasing attention recently [1–4].

The rapid development of visual analytics techniques for machine learning yields an emerging need for a comprehensive review of this area to support the understanding of how visualization techniques are designed and applied to machine learning pipelines. There have been several initial efforts to summarize the advances in this field from different viewpoints. For example, Liu et al. [5] summarized visualization techniques for text analysis. Lu et al. [6] surveyed visual analytics techniques for predictive models. Recently, Liu et al. [1] presented a paper on the analysis of machine learning models from the visual analytics viewpoint. Sacha et al. [7] 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 [5], predictive models [6], model understanding [1]) or aim to sketch an ontology [7] 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 works in the visualization community. Nevertheless, the AI community has also made solid contributions to the study of visually explaining feature detectors in deep learning models. For example, Selvaraju et al. [8] tried to identify the part of an image to which its classification result is sensitive, by computing class

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

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

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

Manuscript received: 2020-07-12; accepted: 2020-08-04

activation maps. Readers can refer to the surveys of Zhang and Zhu [9] and Hohman et al. [3] for more details. We have collected 259 papers from related top-tier venues in the past ten years through a systematical procedure. Based on the machine learning pipeline, we divide this literature as relevant to three stages: before, during, 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, diagnosis, and steering during model building, and data understanding after model building. Each task is illustrated 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 promotes discussion of machine learning related visual analytics techniques and acts as a starting point for practitioners and researchers wishing 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 to develop explainable, trustworthy, and reliable machine learning applications. To comprehensively survey visual analytics techniques for machine learning, we performed an exhaustive manual review of relevant top-tier venues in the past ten years (2010–2020): these were 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 used in a text visualization survey [5]. Specifically, we first considered the titles of papers from these venues to identify candidate papers. Next, we reviewed the abstracts of the candidate papers to further determine whether they concerned visual analytics techniques for machine learning. If the title and abstract did not provide clear information, the full text was gone through to make a final decision. In addition to the exhaustive manual review of the above venues, we also searched for the representative related works that appeared earlier or in other venues, such as the Profiler [10].

After this process, 259 papers were selected. Table 1 presents detailed statistics. Due to the increase in machine learning techniques over the past ten years, this field has been attracting ever more research attention.

2.2 Taxonomy

In this section, we comprehensively analyze the collected visual analytics works to systematically understand the major research trends. These works are categorized based on a typical machine learning pipeline [11] used 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, visual analytics techniques for machine learning can be mapped into these three stages: techniques before model building, techniques during 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 used. Accordingly, there are two research directions, visual analytics for data quality improvement and feature engineering.

Data quality can be improved in various ways, such as completing missing data attributes and correcting wrong data labels. Previously, these tasks were mainly conducted manually or by automatic methods, such as learning-from-crowds algorithms [12] which aim to estimate ground-truth labels from noisy crowd-sourced labels. To reduce experts' efforts or further improve the results of automatic methods, some works employ visual analytics techniques to interactively improve the data quality. Table 1 shows that in recent years, this topic has gained increasing 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 Oriented Gradient) features instead of using raw image pixels. In visual analytics, interactive feature selection provides an interactive and iterative feature selection process. In recent years, in the deep learning era, feature selection and construction are mostly conducted via neural networks. Echoing this trend, there is reducing research attention in recent years (2016–2020) in this direction (see Table 1).

Table 1 Categories of visual analytics techniques for machine learning and representative works in each category (number of papers given in brackets)

Technique category		Papers	Trend
Before model building	Improving data quality (31)	[14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25],[26], [27], [10], [28], [29], [30], [31], [32], [33], [34], [35],[36], [37], [38], [39], [40], [41], [42], [43]	
	Improving feature quality (6)	[44], [45], [46], [47], [48], [49]	
During model building	Model understanding (30)	[50], [51], [52], [53], [54], [55], [56], [57], [58], [59], [60], [61], [62], [63], [64], [65], [66], [67], [68], [69], [70], [71], [72], [73], [74], [75], [76], [77], [78], [79]	
	Model diagnosis (19)	[80], [81], [82], [83], [84], [85], [86], [87], [88], [89], [90], [91], [92], [93], [94], [95], [96], [97], [98]	
	Model steering (29)	[99], [100], [101], [102], [13], [103], [104], [105], [106], [107], [108], [109], [110], [111], [112], [113], [114], [115], [116], [117], [118], [119], [120], [121], [122], [123], [124], [125], [126]	
After model building	Understanding static data analysis results (43)	[127], [128], [129], [130], [131], [132], [133], [134], [135], [136], [137], [138], [139], [140], [141], [142], [143], [144], [145], [146], [147], [148], [149], [150], [151], [152], [153], [154], [155], [156], [157], [158], [159], [160], [161], [162],[163], [164], [165], [166], [167], [168], [169]	
	Understanding dynamic data analysis results (101)	[170], [171], [172], [173], [174], [175], [176], [177], [178], [179], [180], [181], [182], [183], [184], [185], [186], [187], [188], [189], [190], [191], [192], [193], [194], [195], [196], [197], [198], [199], [200], [201], [202], [203], [204], [205], [206], [207], [208], [209], [210], [211], [212], [213], [214], [215], [216], [217], [218], [219], [220], [221], [222], [223], [224], [225], [226], [227], [228], [229], [230], [231], [232], [233], [234], [235], [236], [237], [238], [239], [240], [241], [242], [243], [244], [245], [246], [247], [248], [249], [250], [251], [252], [253], [254], [255], [256], [257], [258], [259], [260], [261], [262], [263], [264], [265], [266], [267], [268], [269], [270]	

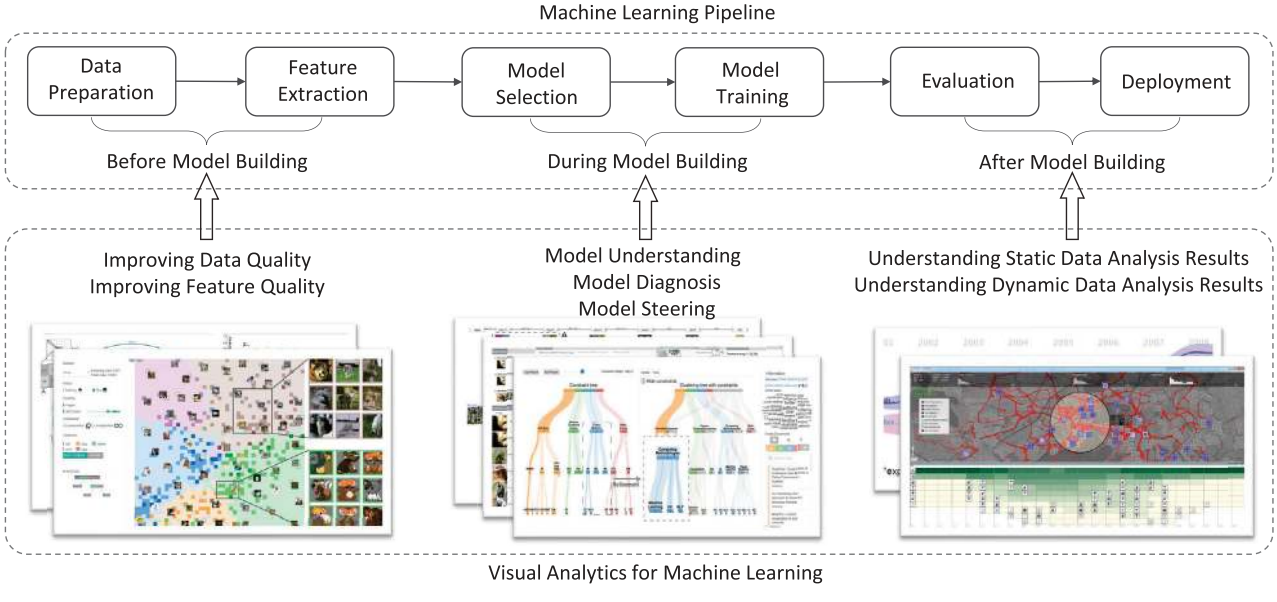

Fig. 1 An overview of visual analytics research for machine learning.

2.2.2 Techniques during 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 (see Table 1). In this survey, we categorize current

methods by their analysis goal: model understanding, diagnosis, or steering. Model understanding methods aim to visually explain the working mechanisms of a model, such as how changes in parameters influence the model and why the model gives a certain output for a specific input. Model diagnosis methods target diagnosing errors in model training via interactive exploration of the training process. Model steering methods are mainly aimed at interactively improving model performance. For example, to refine a topic model, Utopian [13] enables users to interactively merge or split topics, and automatically modify other topics accordingly.

2.2.3 Techniques after model building

After a machine learning model has been built and deployed, it is crucial to help users (e.g., domain experts) understand the model output in an intuitive way, to promote trust in the model output. To this end, there are many visual analytics methods to explore model output, for a variety of applications. Unlike methods for model understanding during model building, these methods usually target model users rather than 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, particularly as static data or temporal data.

3 Techniques before model building

Two major tasks required before building a model are data processing and feature engineering. They are critical, as practical experience indicates that low-quality data and features degrade the performance of machine learning models [271, 272]. Data quality issues include missing values, outliers, and noise in instances and their labels. Feature quality issues include irrelevant features, redundancy between features, etc. While manually addressing these issues is time-consuming, automatic methods may suffer from poor performance. Thus, various visual analytics techniques have been developed to reduce experts' effort as well as to simultaneously improve the performance of automatic methods of producing high-quality data and features [303].

3.1 Improving data quality

Data includes instances and their labels [273]. From this perspective, existing efforts for improving data quality either concern instance-level improvement, or label-level improvement.

3.1.1 Instance-level improvement

At the instance level, many visual analytics methods focus on detecting and correcting anomalies in data, such as missing values and duplication. For example, Kandel et al. [10] 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. Then, linked summary visualizations are automatically recommended to facilitate the discovery of potential causes and consequences of these anomalies. VIVID [14] was developed to handle missing values in longitudinal cohort study data. Through multiple coordinated visualizations, experts can identify the root causes of missing values (e.g., a particular group who do not participate in follow-up examinations), and replace missing data using an appropriate imputation model. Anomaly removal is often an iterative process which must be repeated. Illustrating provenance in this iterative process allows users to be aware of changes in data quality and to build trust in the processed data. Thus, Bors et al. [20] proposed DQProv Explorer to support the analysis of data provenance, using a provenance graph to support the navigation of data states and a quality flow to present changes in data quality over time. Recently, another type of data anomaly, out-of-distribution (OoD) samples, has received extensive attention [274, 275]. 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. [21] proposed OoDAnalyzer to detect and analyze OoD samples. An ensemble OoD detection method, combining both high- and low-level features, was proposed to improve detection accuracy. A grid visualization of the detection result (see Fig. 2) is utilized to explore OoD samples in context and explain the underlying reasons for their presence. In order to generate grid layouts at 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

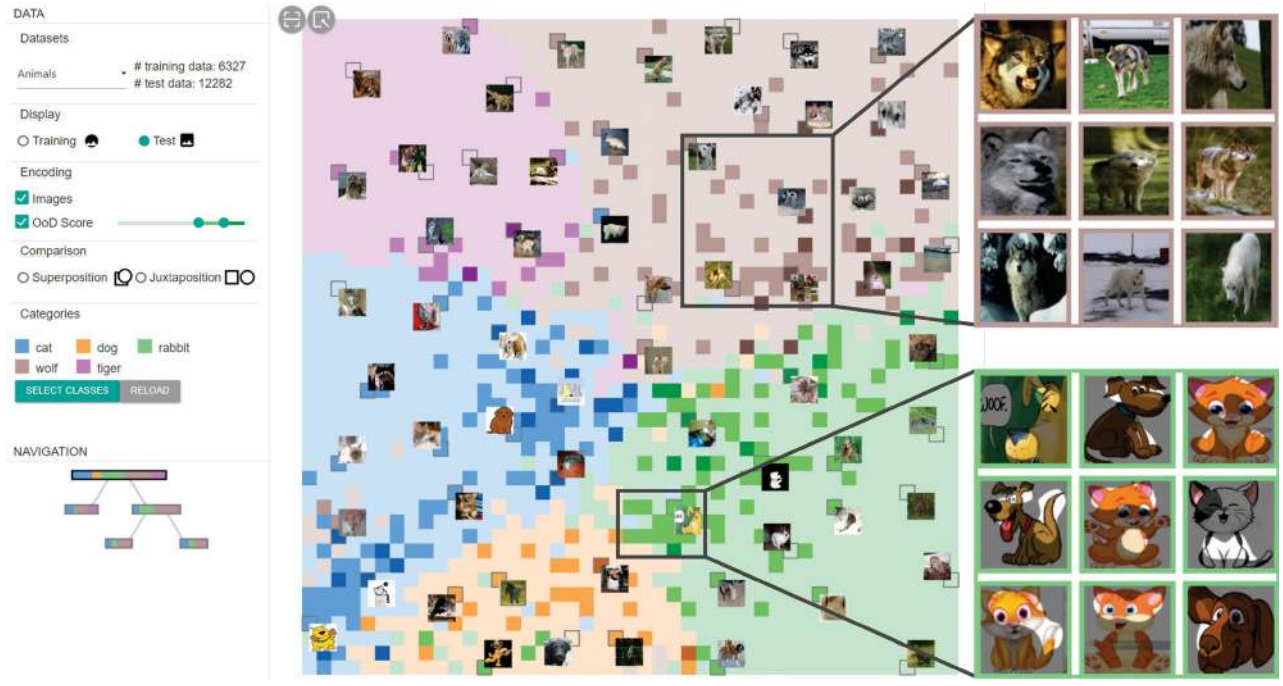

Fig. 2 OoDAnalyzer, an interactive method to detect out-of-distribution samples and explain them in context. Reproduced with permission from Ref. [21], © IEEE 2020.

that induce specific quality issues that require analysis in a temporal context. To tackle this issue, Arbesser et al. [15] proposed a visual analytics system, Visplause, to visually assess time-series data quality. 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 analysis of a group of anomalies (e.g., abnormal time series of the same type) simultaneously. Besides automatically detected anomalies, KYE [23] also supports the identification of additional anomalies overlooked by automatic methods. Time-series data are presented in a heatmap view, where abnormal patterns (e.g., regions with unusually high values) indicate potential anomalies. Click stream data are a widely studied kind of time-series data in the field of visual analytics. To better analyze and refine click stream data, Segmentifier [22] was proposed to provide an iterative exploration process for segmentation and analysis. Users can explore segments in three coordinated views at different granularities and refine them by filtering, partitioning, and transformation. Every refinement step results in new segments, which can be further analyzed and refined.

To tackle uncertainties in data quality improvement, Bernard et al. [17] developed a visual analytics tool to exhibit the changes in the data and uncertainties caused by different preprocessing methods. This tool enables experts to become aware of the effects of these methods and to choose suitable ones, to reduce task-irrelevant parts while preserving task-relevant parts of the data.

As data have the risk of exposing sensitive information, several recent studies have focused on preserving data privacy during the data quality improvement process. For tabular data, Wang et al. [41] 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 changes as privacy-preserving operations are applied. To preserve privacy in network datasets, Wang et al. [40] 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 low priorities, reducing the possibility of modifying these nodes. Based on node priorities and utility metrics, users can apply and compare a set of privacy-preserving operations and choose the most suitable one according to their knowledge and experience.

3.1.2 Label-level improvement

According to whether the data have noisy labels, existing works can be classified as either methods for improving the quality of noisy labels, or allowing interactive labeling.

Crowdsourcing provides a cost-effective way to collect labels. However, annotations provided by crowd workers are usually noisy [271, 276]. Many methods have been proposed to remove noise in labels. Willett et al. [42] developed a crowd-assisted clustering method to remove redundant explanations provided by crowd workers. Explanations are clustered into groups, and the most representative ones are preserved. Park et al. [35] proposed C^2A that visualizes crowdsourced annotations and worker behavior to help doctors identify malignant tumors in clinical videos. Using C^2A, doctors can discard most tumor-free video segments and focus on the ones that most likely to contain tumors. To analyze the accuracy of crowdsourcing workers, Park et al. [34] developed CMed that visualizes clinical image annotations by crowdsourcing, and workers' behavior. By clustering workers according to their annotation accuracy and analyzing their logged events, experts are able to find good workers and observe the effects of workers' behavior patterns. LabelInspect [31] was proposed to improve crowdsourced labels by validating uncertain instance labels and unreliable workers. Three coordinated visualizations, a confusion (see Fig. 3(a)), an instance (see Fig. 3(b)), and a worker visualization (see Fig. 3(c)), were developed to facilitate the identification and validation of uncertain instance labels and unreliable workers. Based on expert validation, further instances and workers are recommended for validation by an iterative and progressive verification procedure.

Although the aforementioned methods can effectively improve crowdsourced labels, crowd information is not available in many real-world datasets. For example, the ImageNet dataset [277] only contains the labels cleaned by automatic noise removal methods. To tackle these datasets, Xiang et al. [43] developed DataDebugger to interactively improve data quality by utilizing user-selected trusted items. Hierarchical visualization combined with an incremental projection method and an outlier biased sampling method facilitates the exploration and identification of trusted items. Based on these identified trusted items, a data correction algorithm propagates labels from trusted items to the whole dataset. Paiva et al. [33] assumed that instances misclassified by a trained classifier were likely to

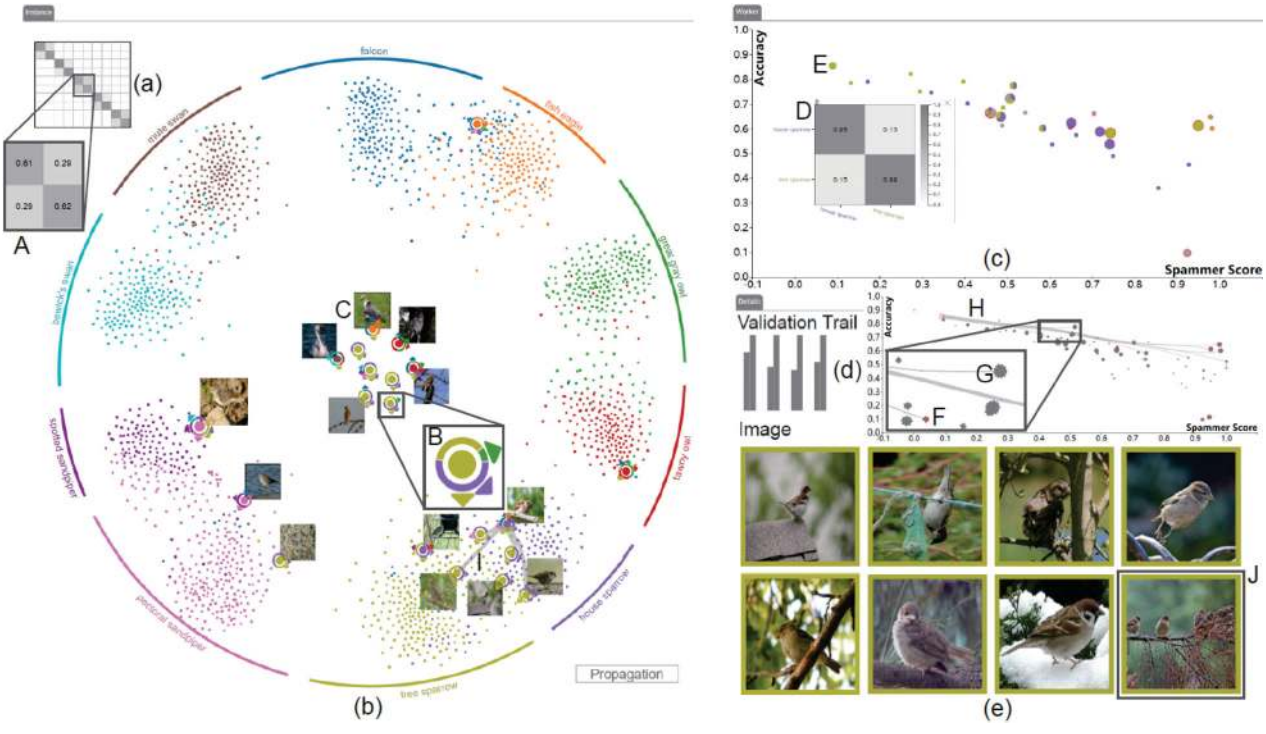

Fig. 3 LabelInspect, an interactive method to verify uncertain instance labels and unreliable workers. Reproduced with permission from Ref. [31], © IEEE 2019.

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 correction, the classifier is refined using the corrected labels, and a new round of correction starts. Bäuerle et al. [16] developed three classifier-guided measures to detect data errors. Data errors are then presented in a matrix and a scatter plot, allowing 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 effort is a major goal of interactive labeling. To this end, Moehrmann et al. [32] used an SOM-based visualization to place similar images together, allowing users to label multiple similar images of the same class in one go. This strategy is also used by Khayat et al. [28] to identify social spambot groups with similar anomalous behavior, Kurzhals et al. [29] to label mobile eye-tracking data, and Halter et al. [24] to annotate and analyze primary color strategies used in films. Apart from placing similar items together, other strategies, like filtering, have also been applied to find items of interest for labeling. Filtering and sorting are utilized in MediaTable [36] to find similar video segments. A table visualization is utilized to present video segments and their attributes. Users can filter out irrelevant segments and sort on attributes to order relevant segments, allowing users to label several segments of the same class simultaneously. Stein et al. [39] provided a rule-based filtering engine to find 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 have combined visualization techniques with machine learning techniques, such as active learning. The concept of "intra-active labeling" was first introduced by Höferlin et al. [26]; it 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 [25], sequential data retrieval [30], trajectory classification [27], identifying relevant tweets [37], and argumentation mining [38]. For example, to annotate text fragments in argumentation mining tasks, Sperrle et al. [38] developed a language model for fragment recommendation. A layered visual abstraction is utilized to support five relevant analysis tasks required by text fragment annotation. In addition to developing systems for interactive labeling, some empirical experiments were conducted to demonstrate their effectiveness. For example, Bernard et al. [18] conducted experiments to show the superiority of user-centered visual interactive labeling over model-centered active learning. A quantitative analysis [19] was also performed to evaluate user strategies for selecting samples in the labeling process. Results show that in early phases, data-based (e.g., clusters and dense areas) user strategies work well. However, in later phases, model-based (e.g., class separation) user strategies 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 [278]. A common feature selection strategy is to select a subset of features that minimizes the redundancy between them and maximizes the relevance between them and targets (e.g., classes of instances) [46]. Along this line, several methods have been developed to interactively analyze the redundancy and relevance of features. For example, Seo and Shneiderman [48] proposed a rank-by-feature framework, which ranks features by relevance. They visualized ranking results with tables and matrices. Ingram et al. [44] proposed a visual analytics system, DimStiller, which allows users to explore features and their relationships and interactively remove irrelevant and redundant features. May et al. [46] proposed SmartStripes to select different feature subsets for different data subsets. A matrix-based layout is utilized to exhibit the relevance and redundancy of features. Mühlbacher and Piringer [47] 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 detail. Parallel coordinates visualization was utilized by Tam et al. [49] to identify features that could discriminate between different

clusters. Krause et al. [45] 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 in model building. For example, FeatureInsight [279] was proposed to construct new features for text classification. By visually examining classifier errors and summarizing the root causes of these errors, users are able to create new features that can correctly discriminate misclassified documents. To improve the generalization capability of new features, visual summaries are used to analyze a set of errors instead of individual errors.

4 Techniques during 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 how to reveal the underlying working mechanisms of machine learning models and then help model developers to build well-formed models. First of all, model developers require a comprehensive understanding of models in order to release them from a time-consuming trial-and-error process. When the training process fails or the model does not provide satisfactory performance, model developers need to diagnose the issues occurring in the training process. Finally, there is a need to assist in model steering as much time is spent in improving model performance during the model building process. Echoing these needs, researchers have developed many visual analytics methods to enhance model understanding, diagnosis, and steering [1, 2].

4.1 Model understanding

Works related to model understanding belong to two classes: those understanding the effects of parameters, and those understanding model behaviour.

4.1.1 Understanding the effects of parameters

One aspect of model understanding is to inspect how the model outputs change with changes in model parameters. For example, Ferreira et al. [54] developed BirdVis to explore the relationships between different parameter configurations and model outputs; these were 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. [266] proposed a visual analytics method to visualize how variables affect statistical indicators in a logistic regression model.

4.1.2 Understanding model behaviours

Another aspect is how the model works to produce the desired outputs. There are three main types of methods used to explain model behaviours, namely network-centric, instance-centric, and hybrid methods. Network-centric methods aim to explore the model structure and interpret how different parts of the model (e.g., neurons or 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 [280], but visual clutter becomes a serious problem as the model structure becoming increasingly complex. To tackle this problem, Liu et al. [62] developed CNNVis to visualize deep convolutional neural networks (see Fig. 4). It leverages clustering techniques to group neurons with similar roles as well as their connections in order to address visual clutter caused by their huge quantity. This tool helps experts understand the roles of the neurons and their learned features, and moreover, how low-level features are aggregated into high-level ones through the network. Later, Wongsuphasawat et al. [77] designed a graph visualization for exploring the machine learning model architecture in Tensorflow [281]. 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 to provide instance-level analysis and exploration, as well as understanding the relationships between instances. Rauber et al. [69] 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 evolves during training so as to understand the network's learning behaviour. Some visual analytics techniques for understanding recurrent neural

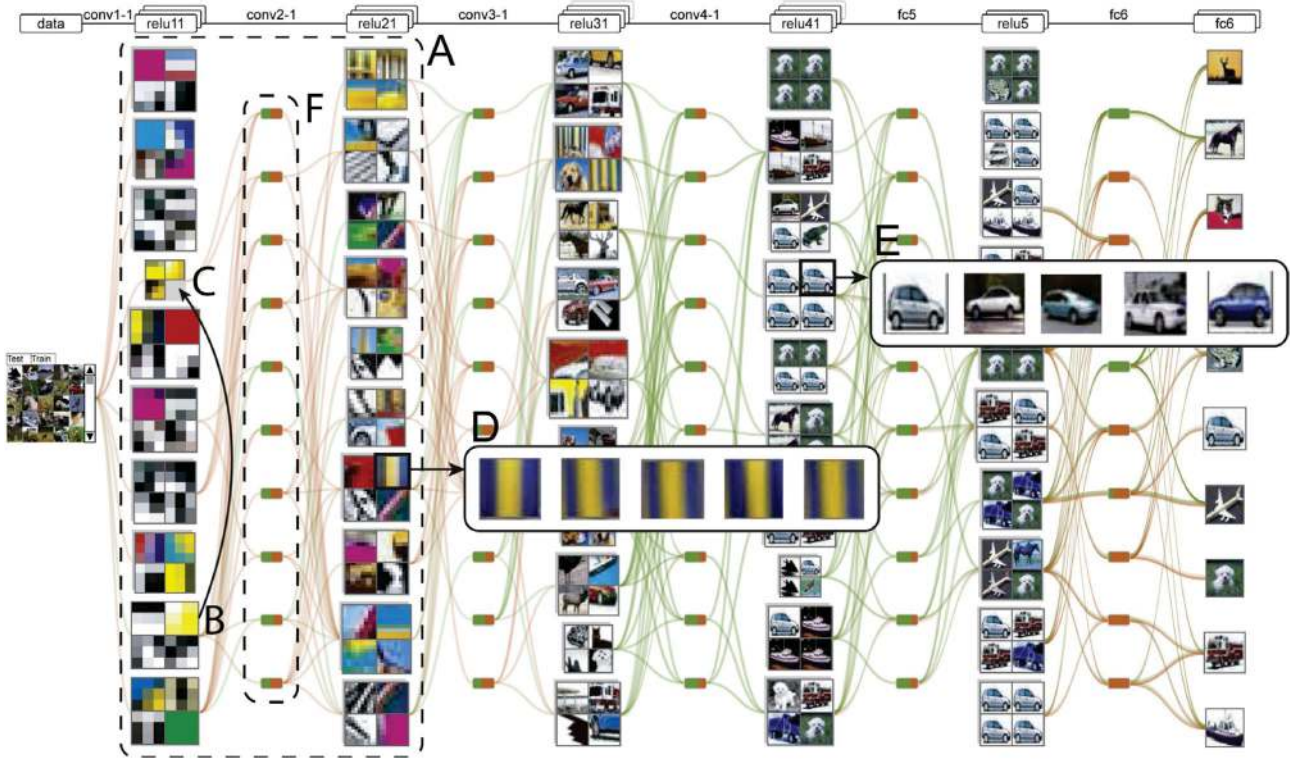

Fig. 4 CNNVis, a network-centric visual analytics technique to understand deep convolutional neural networks with millions of neurons and connections. Reproduced with permission from Ref. [62], © IEEE 2017.

networks (RNNs) also adopt such an instance-centric design. LSTMVis [73] developed by Strobelt et al. utilizes parallel coordinates to present the hidden states, to support the analysis of changes in the hidden states over texts. RNNVis [65] developed by Ming et al. clusters the hidden state units (each hidden state unit is a dimension of the hidden state vector in an RNN) as memory chips and words as word clouds. Their relationships are modeled as a bipartite graph, which supports sentence-level explanations in RNNs.

Hybrid methods combine the above two methods and leverage both of their strengths. In particular, instance-level analysis can be enhanced with the context of the network architecture. Such contexts benefit the understanding of the network's working mechanism. For instance, Hohman et al. [56] proposed Summit, to reveal important neurons and critical neuron associations contributing to the model prediction. It integrates an embedding view to summarize the activations between classes and an attribute graph view to reveal influential connections between neurons. Kahng et al. [59] proposed ActiVis for large-scale deep neural networks. It visualizes the model structure with a computational graph and the activation relationships between instances, subsets, and classes using a projected view.

In recent years, there have been some efforts to use a surrogate explainable model to explain model behaviours. The major benefit of such methods is that they do not require users to investigate the model itself. Thus, they are more useful for those with no or limited machine learning knowledge. Treating the classifier as a black box, Ming et al. [66] first extracted rule-based knowledge from the input and output of the classifier. These rules are then visualized using RuleMatrix, which supports interactive exploration of the extracted rules by practitioners, improving the interpretability of the model. Wang et al. [75] developed DeepVID to generate a visual interpretation for image classifiers. Given an image of interest, a deep generative model was first used to generate 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.

4.2 Model diagnosis

Visual analytical techniques for model diagnosis may either analyze the training results, or analyze the training dynamics.

4.2.1 Analyzing training results

Tools have been developed for diagnosing classifiers based on their performance [81, 82, 86, 93]. For example, Squares [93] used boxes to represent samples and group them according to their prediction classes. Using different textures to encode true/false positives/negatives, this tool allows fast and accurate estimation of performance metrics at multiple levels of detail. Recently, the issue of model fairness has drawn growing attention [80, 83, 97]. For example, Ahn et al. [80] 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 individual and group levels using different measures. Based on these measures, developers can iteratively identify those features that cause discrimination and remove them from the model. Researchers are also interested in exploring potential vulnerabilities in models that prevent them from being reliably applied to real-world applications [84, 91]. Cao et al. [84] proposed AEVis to analyze how adversarial examples fool neural networks. The system (see Fig. 5) takes both normal and adversarial examples as input and extracts their datapaths for model prediction. It then employs a river-based metaphor to show the diverging and merging patterns of the extracted datapaths, which reveal where the adversarial samples mislead the model. Ma et al. [91] designed a series of visual representations from overview to detail to reveal how data poisoning will make a model misclassify a specific sample. By comparing the distributions of the poisoned and normal training data, experts can deduce the reason for the misclassification of the attacked sample.

4.3 Analyzing training dynamics

Recent efforts have also been concentrated on analyzing the training dynamics. These techniques are intended for debugging the training process of machine learning models. For example, DGMTracker [89] assists experts to discover reasons for 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

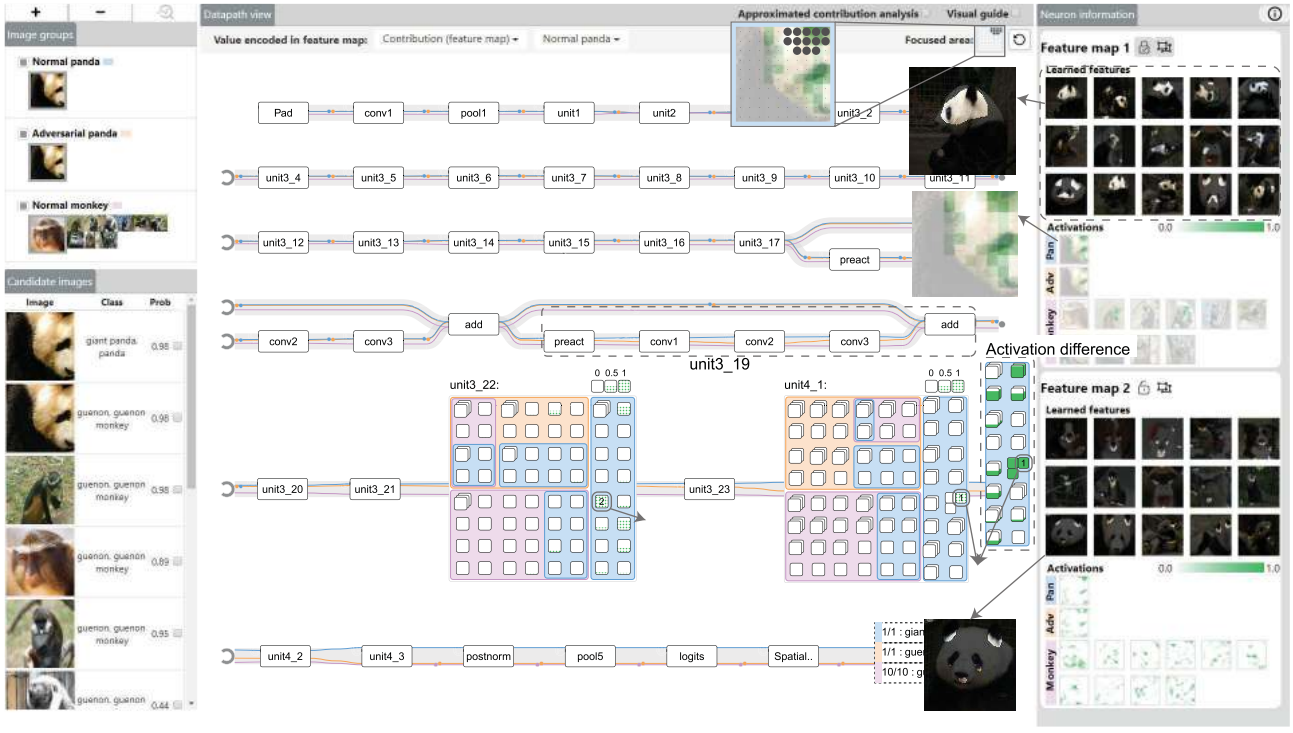

Fig. 5 AEVis, a visual analytics system for analyzing adversarial samples. It shows diverging and merging patterns in the extracted datapaths with a river-based visualization, and critical feature maps with a layer-level visualization. Reproduced with permission from Ref. [84], © IEEE 2020.

experts detect the potential root cause of a failure. It also employs a credit assignment algorithm to disclose the interactions between neurons to facilitate the diagnosis of failure propagation. Attention has also been given to the diagnosis of the training process of deep reinforcement learning. Wang et al. [96] proposed DQNViz for the understanding and diagnosis of deep Q-networks for a Breakout game. At the overview level, DQNViz presents changes in the overall statistics during the training process with line charts and stacked area charts. Then at the detail level, it uses segment clustering and a pattern mining algorithm to help experts identify common as well as suspicious patterns in the event-sequences of the agents in Q-networks. As another example, He et al. [87] proposed DynamicsExplorer to diagnose an LSTM trained to control a ball-in-maze game. To support quick identification of where training failures arise, it visualizes ball trajectories with a trajectory variability plot, as well as their clusters using a parallel coordinates plot.

4.4 Model steering

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

4.4.1 Model refinement with human knowledge

Several visual analytics techniques have been developed to place users into the loop of the model refinement process, through flexible interaction.

Users can directly refine the target model with visual analytics techniques. A typical example is ProtoSteer [116], a visual analytics system that enables editing prototypes to refine a prototype sequence network named ProSeNet [282]. ProtoSteer uses four coordinated views to present the information about the learned prototypes in ProSeNet. Users can refine these prototypes by adding, deleting, and revising specific prototypes. The model is then retrained with these user-specific prototypes for performance gain. In addition, van der Elzen and van Wijk [122] proposed BaobabView to support experts to construct decision trees iteratively using domain knowledge. Experts can refine the decision tree with direct operations, including growing, pruning, and optimizing the internal nodes, and can evaluate the refined one with various visual representations.

Besides direct model updates, users can also correct flaws in the results or provide extra knowledge, allowing the model to 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 [13, 105, 106, 109, 124, 125]. For instance, Yang et al. [125] presented ReVision that allows users to steer 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

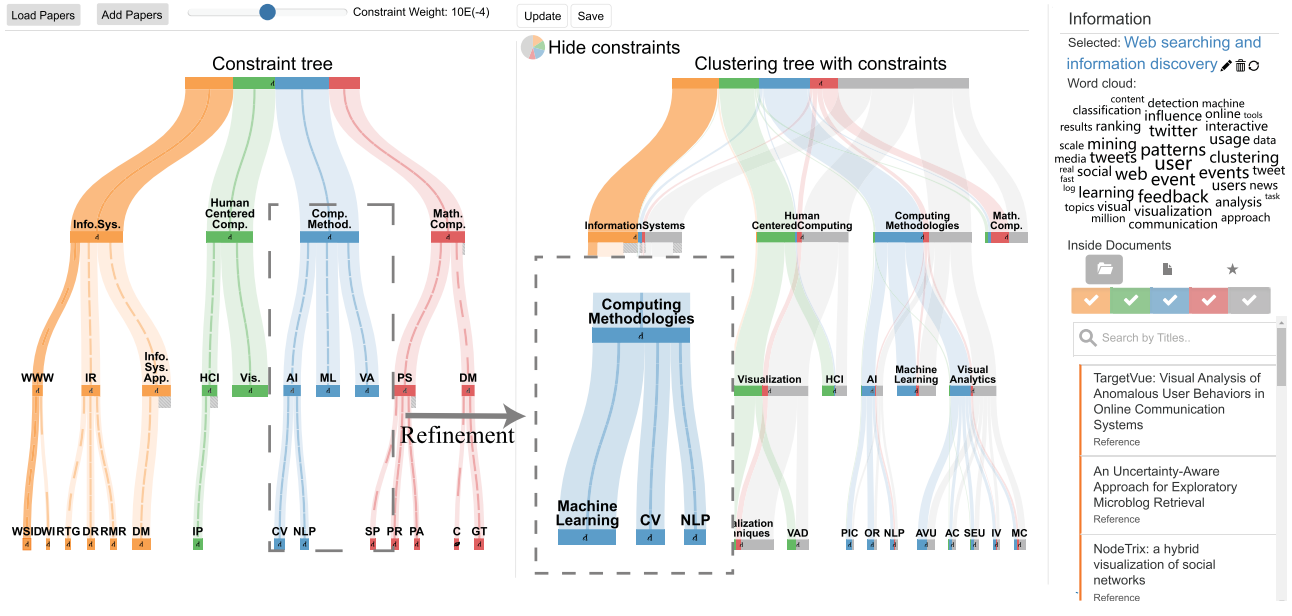

Fig. 6 ReVision, a visual analytics system integrating a constrained hierarchical clustering algorithm with an uncertainty-aware, tree-based visualization to help users interactively refine hierarchical topic modeling results. Reproduced with permission from Ref. [125], © IEEE 2020.

uncertainty-aware tree-based visualization to guide the steering of the clustering results. Users can refine the constraint hierarchy by dragging. Documents are then re-clustered based on the modified constraints. Other human-in-the-loop models have also stimulated the development of visual analytic systems to support such kinds of model refinement. For instance, Liu et al. [112] proposed MutualRanker using an uncertainty-based mutual reinforcement graph model to retrieve important blogs, users, and hashtags from microblog data. It shows ranking results, uncertainty, and its propagation with the help of a composite visualization; users can examine the most uncertain items in the graph and adjust their ranking scores. The model is incrementally updated by propagating adjustments throughout the graph.

4.4.2 Model selection from an ensemble

Another strategy for model steering is to select the best model from a model ensemble, which is usually found in clustering [102, 118, 121] and regression models [99, 103, 113, 119]. Clustrophile 2 [102] is a visual analytics system for visual clustering analysis, which guides user selection of appropriate input features and clustering parameters through recommendations based on user-selected results. BEAMES [103] was designed for multimodel steering and selection in regression tasks. It creates a collection of regression models by varying algorithms and their corresponding hyperparameters, with further optimization by 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 to help users understand and gain insights from model outputs, such as high-dimensional data analysis results [5, 283]. As these methods are often data-driven, we categorize the corresponding methods according to the type of data analyzed. The temporal property of data is critical in visual design. Thus, we classify methods as those understanding static data analysis results, and those understanding dynamic data analysis results. A visual analytics system for understanding static data analysis results usually treats all model output as a large collection and analyzes the static structure. For dynamic data, in addition to understanding the analysis results at each time point, the system focuses 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 the type of data. Most research focuses on textual data analysis, while fewer works study the understanding of other types of data analysis.

5.1.1 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 a large amount of textual data [5].

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. [143] 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 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 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 provide deeper insights into textual data. To provide an overview of the relevant topics discussed in multiple sources, Liu et al. [159] first utilized a correlated topic model to extract topic graphs from multiple text sources. A graph matching algorithm is then developed to match the topic graphs from different sources, and a hierarchical clustering method is employed to generate hierarchies of topic graphs. Both the matched topic graph and hierarchies are fed into a hybrid visualization which consists of a radial icicle plot and a density-based node-link diagram (see Fig. 7(a)), to support exploration and analysis of common and distinctive topics discussed

in multiple sources. Dou et al. [136] 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) [284], is employed to extract topic features from the corpus. Relationships between the demographic information and extracted features are explored through a parallel sets visualization [285], and different demographic groups are projected onto the two-dimensional space based on the similarity of their topics of interest (see Fig. 7(b)). Recently, some deep learning models have also been adopted because of their better performance. For example, Berger et al. [128] proposed cite2vec to visualize the latent themes in a document collection via document usage (e.g., citations). It extended a famous word2vec model, the skip-gram model [286], to generate the embedding for both words and documents by considering the citation information and the textual content together. The words are projected into a two-dimensional 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.

5.1.2 Other data analysis

In addition to textual data, other types of data have also been studied. For example, Hong et al. [146] analyzed flow fields through an LDA model by defining pathlines as documents and features as words, respectively. After modeling, the original pathlines and extracted topics were projected into a two-dimensional space using multidimensional scaling, and several previews were generated to render the pathlines for important topics. Recently, a visual analytics tool, SMARTexplore [129], was developed to help analysts find and understand interesting patterns within and between 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 latent themes in data change over time. For example, a system can help politicians to make timely decisions if it provides an overview of major public opinions on social media and how they change over time. Most 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 may further classify existing works on visual dynamic data analysis as

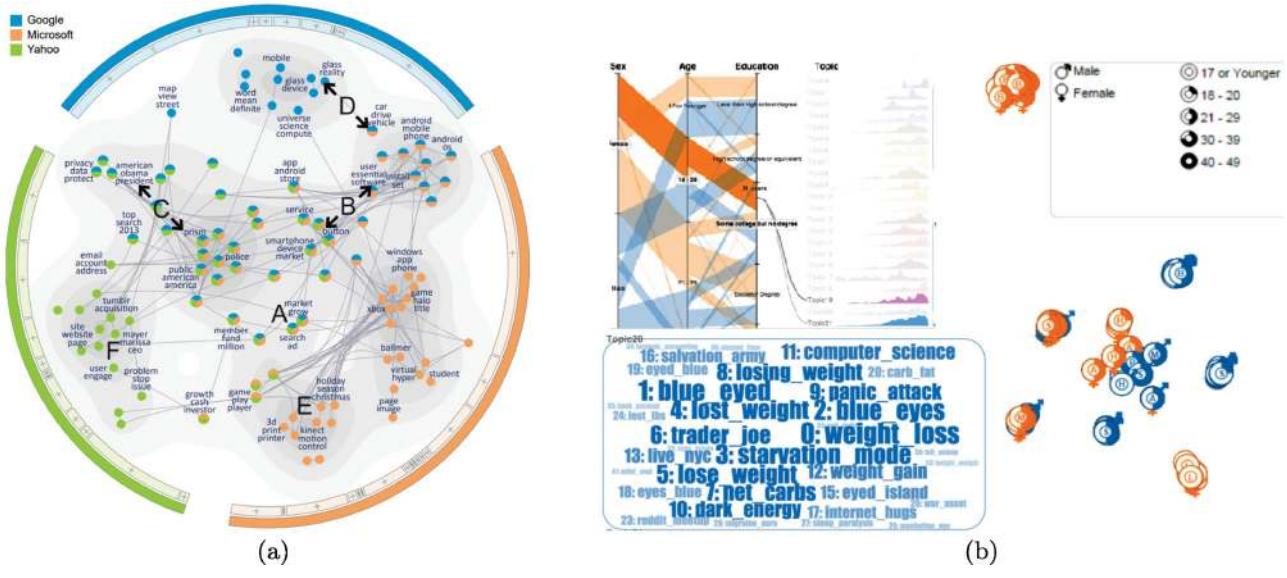

Fig. 7 Examples of static text visualization. (a) TopicPanorama extracts topic graphs from multiple sources and reveals relationships between them using graph layout. Reproduced with permission from Ref. [159], © IEEE 2014. (b) DemographicVis measures similarity between different users after analyzing their posting contents, and reveals their relationships using t-SNE projection. Reproduced with permission from Ref. [136], © IEEE 2015.

offline and online. In offline analysis, all data are available before analysis, while online analysis tackles streaming data that is incoming during the analysis process.

5.2.1 Offline analysis.

Offline analysis research can be classified according to the analysis task: topic analysis, event analysis, and trajectory analysis.

Understanding topic evolution in a large text corpus over time is an important topic, attracting much attention. Most existing works adopt a river metaphor to convey changes in the text corpus over time. ThemeRiver [204] is one of the pioneering works, using the river metaphor to reveal changes in the volumes of different themes. To better understand the content change of a document corpus, TIARA [220, 248] utilizes an LDA model [287] to extract topics from the corpus and reveal their changes over time. However, only observing volumes and content change is not enough for complex analysis tasks where users want to explore relationships between different topics and their changes over time. Therefore, later works have focused on understanding relationships between topics (e.g., topic splitting and merging) and their evolving patterns over time. For example, Cui et al. [190] first extracted topic splitting and merging patterns from a document collection using an incremental hierarchical Dirichlet process model [288]. 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. [259] leveraged a topic competition model to extract dynamic competition between topics and the effects of opinion leaders on social media. Sun et al. [238] extended the competition model to a "coopetition" (cooperation and competition) model to help understand the more complex interactions between evolving topics. Wang et al. [246] proposed IdeaFlow, a visual analytics system for learning the lead-lag relationships across different social groups over time. However, these works use a flat structure to model 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. [191] extracted a sequence of topic trees using an evolutionary Bayesian rose tree algorithm [289] and then calculated the tree cut for each tree. These tree cuts are used to approximate the topic trees and display them in a river metaphor, which also reveals dynamic relationships between the topics, including topic birth, death, splitting, and merging.

Event analysis targets revealing common or semantically important sequential patterns in ordered sequences of events [149, 202, 222, 226]. To facilitate visual exploration of large scale event sequences and pattern discovery, several visual analytics methods have been proposed. For example, Liu et al. [222] developed a visual analytics method for click stream data. Maximal sequential patterns are discovered and

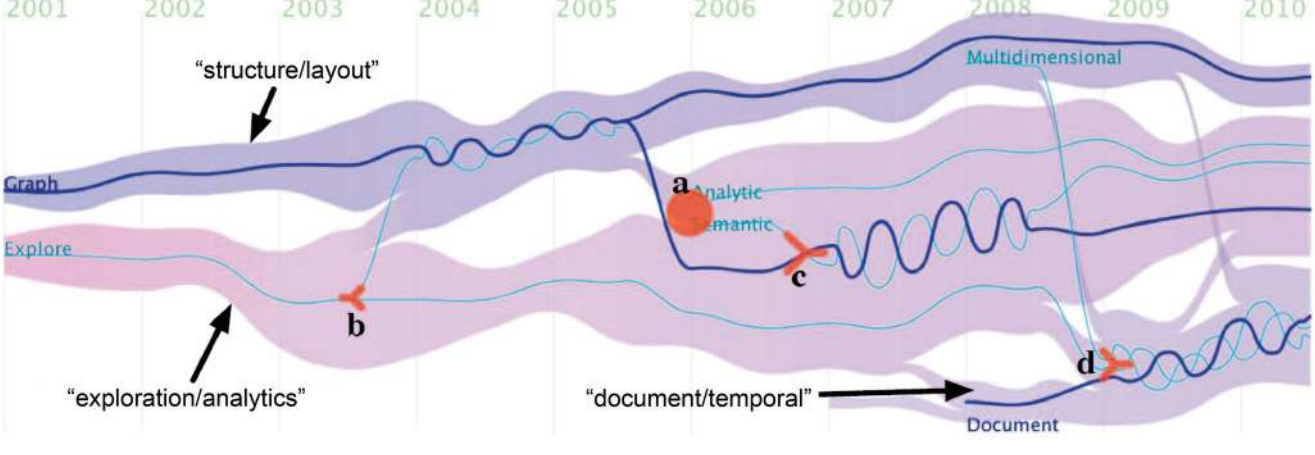

Fig. 8 TextFlow employs a river-based metaphor to show topic birth, death, merging, and splitting. Reproduced with permission from Ref. [190], © IEEE 2011.

pruned from the click stream data. The extracted patterns and original data are well illustrated at four granularities: patterns, segments, sequences, and events. Guo et al. [202] developed EventThread, which uses a tensor-based model to transform the event sequence data into an n-dimensional tensor. 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 [201]. 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 was developed to reveal and analyze the evolution patterns across stages.

Other works focus on understanding movement data (e.g., GPS records) analysis results. Andrienko et al. [174] extracted movement events from trajectories and then performed spatio-temporal clustering for aggregation. These clusters are visualized using spatio-temporal envelopes to help analysts find potential traffic jams in the city. Chu et al. [189] adopted an LDA model for mining latent movement patterns in taxi trajectories. The movement of each taxi, represented by the traversed street names, was regarded as a document. Parallel coordinates were 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. [269] treated 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 into two-dimensional space, where analysts can check the distributions of the origin-destination flows and select some for display on the map. Besides directly analyzing the original trajectory data, other papers try to augment the trajectories with auxiliary information to reduce the burden on visual exploration. Kruger et al. [212] clustered destinations with DBScan and then used Foursquare to provide detailed information about the destinations (e.g., shops, university, residence). Based on the enriched data, frequent patterns were extracted and displayed in the visualization (see Fig. 9); icons on the time axis help understand these patterns. Chen et al. [186] mined trajectories from geo-tagged social media and displayed keywords extracted from the text content, helping users explore the semantics of trajectories.

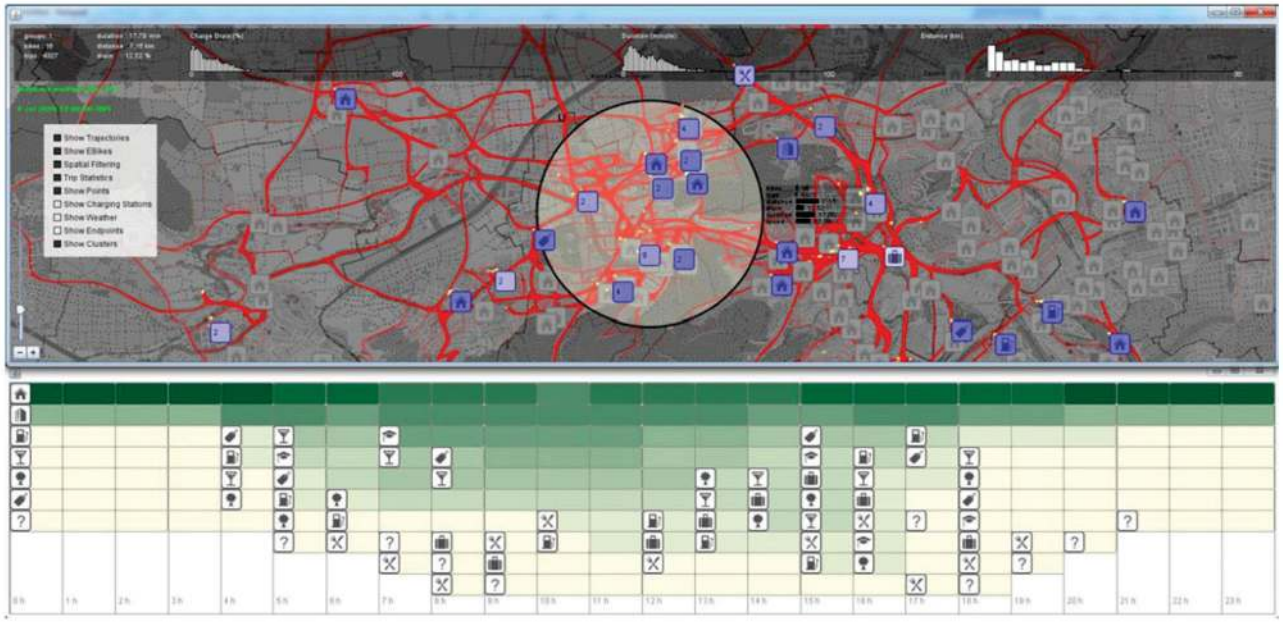

Fig. 9 Kruger et al. enrich trajectory data semantically. Frequent routes and destinations are visualized in the geographic view (top), while the frequent temporal patterns are mined and displayed in the temporal view (bottom). Reproduced with permission from Ref. [212],© IEEE 2015.

5.2.2 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 extracts 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, spatio-temporal anomalies were detected and reported to users in real-time. To reduce user effort for filtering and monitoring in ScatterBlogs, Bosch et al. [177] 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. [219] 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 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 newly arriving textual documents are gradually aggregated into the existing topics over time. Triggered by TopicStream, Wu et al. [252] developed StreamExplorer, which enables the tracking and comparison of a social stream. In particular, an entropy-based event detection method was developed to detect events in the social media stream. They 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. [213] employed a long short-term memory model for road traffic congestion forecasting and visualized the results with a Volume-Speed Rivers visualization. Propagation of congestion was also extracted and visualized, helping analysts understand causality within the detected congestion.

6 Research opportunities

Although visual analytics research for machine learning has achieved promising results in both academia 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

6.1.1 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 [290]. Most existing methods focus on inaccurate data (e.g., noisy crowdsourced annotations and label errors) quality issues, and interactive labeling related to incomplete data (e.g., none or only a few data are labeled) quality issues. However, fewer efforts are devoted to the better exploitation of unlabeled data related to incomplete data quality issues as well as inexact data (e.g., coarse-grained labels that are not exact as required) quality issues. This paves the way for potential future research. Firstly, the potential of visual analytics techniques to address the incompleteness issue is not fully exploited. For example, improving the quality of unlabeled data is critical for semi-supervised learning [290, 291], which is tightly combined with a small amount of labeled data during training to infer the correct mapping from the data set to the label set. One typical example is graph-based semi-supervised learning [291], 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 these poor-quality graphs is that automatic graph construction methods usually rely on global parameters (e.g., a global k value in the kNN graph construction method), which may be locally inappropriate. As a consequence, it is necessary to utilize visualization to illustrate how labels are propagated along graph edges, to facilitate understanding of how local graph structures affect model performance. Based on such understanding, experts can adaptively modify the graph to gradually create a higher-quality graph.

Secondly, although the inexact data quality issue

is common in real-world applications [292], it has received little attention from the field of visual analytics. This issue refers to the situation where labels are inexact, e.g., coarse-grained labels, such as arise in computed tomography (CT) scans. The labels of CT scans usually come from corresponding diagnosis reports that describe whether patients have certain medical problems (e.g., a tumor). For a CT scan with tumors, we only know that one or more slices in the scan contain tumors. However, we do not know which slices contain tumors as well as the exact tumor locations in these slices. Although various machine learning methods [293, 294] have been proposed to learn from such coarse-grained labels, they may lead to poor performance [290] due to the lack of exact information. Fine-grained validation is 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 predictions, to develop an interactive verification process for providing more finely-grained labels while minimizing expert effort.

6.1.2 Explainable feature engineering

Most existing work for improving feature quality focuses on tabular or textual data from traditional analysis models. The features of these data are naturally interpretable, which makes the feature engineering process simple. However, features extracted by deep neural networks perform better than handcrafted ones [295, 296]. These deep features are hard to interpret due to the black box nature of deep neural networks, which brings several challenges for feature engineering.

Firstly, the extracted features are obtained in a data-driven process, which may poorly represent the original images/videos when the datasets are biased. For example, given a dataset with only dark dogs and light cats, the extracted features may emphasize color and ignore other discriminating concepts, like shapes of faces and ears. Without a clear understanding of these biased features, it is hard to correct them in a suitable 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 discarded by the extracted features and to visualize it in a comprehensible manner.

Moreover, redundancy exists in extracted deep features [297]. Removing redundant features can lead to several benefits, such as reducing storage requirements and improving generalization [278]. However, without a clear understanding of the exact meaning of features, it is hard to judge whether a feature is redundant. Thus, an interesting future topic is to develop a visual analytics method to convey feature redundancy in a comprehensible way, to allow experts to explore it, and remove redundant features.

6.2 Opportunities during model building

6.2.1 Online training diagnosis

Existing visual analytics tools for model diagnosis mostly work offline: the data for diagnosis is collected after the training process is finished. They have shown their capability for revealing the root causes of failed training processes. However, as modern machine learning models become more and more complex, training processes can last for days or even weeks. Offline diagnosis severely restricts the ability of visual analytics to assist in training. Thus, there is a significant need to develop visual analytics tools for online diagnosis of the training process so that model developers can identify anomalies and promptly make corresponding adjustments to the process. This can save much time in the trial-and-error model building process. The key challenge for online diagnosis is to detect anomalies in the training process in a timely manner. While it remains a difficult task to develop algorithms for automatically and accurately detecting anomalies in real-time, interactive visualization promises a way to locate potential errors in the training process. Differing from offline diagnosis, 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 partial streaming data. These techniques can help experts monitor online model training processes and identify possible issues rapidly.

6.2.2 Interactive model refinement

Recent works have explored the utilization of uncertainty to facilitate interactive model refinement [106, 112, 124, 125]. There are many methods to assign uncertainty scores to model outputs (e.g., based on confidence scores produced by classifiers),

and visual hints can be used to guide users to examine model outputs with high uncertainty. Models uncertainty will be recomputed after user refinement, and users can perform iteratively until they are satisfied with the results. Furthermore, additional information can also be leveraged to provide users with more intelligent guidance to facilitate a fast and accurate model refinement process. However, the room for improving interactive model refinement is still largely unexplored by researchers. One possible direction is that since the refinement process usually requires several iterations, guidance in later iterations can be learned from users' previous 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, prior knowledge can be used to predict where refinements are needed. For example, model outputs may conflict with certain public or domain knowledge, especially for unsupervised models (e.g., nonlinear matrix factorization and latent Dirichlet allocation for topic modeling), which should be considered in the refinement process. Therefore, such a knowledge-based strategy focuses on revealing unreasonable results produced by the models, allowing users to refine the models by adding constraints to them.

6.3 Opportunities after model building

6.3.1 Understanding multi-modal data

Existing works on content analysis have achieved great success in understanding single-modal data, such as texts, images, and videos. However, real-world applications often contain multi-modal data, which combines several different content forms, such as text, audio, and images. For example, a physician diagnoses a patient after considering multiple kinds of data, such as the medical record (text), laboratory reports (tables), and CT scans (images). When analyzing such multi-modal data, in-depth relationships between different modals cannot be well captured by simply combining knowledge learned from single-modal models. It is more promising to employ multi-modal machine learning techniques and leverage their 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 such multi-modal learning models. Many machine learning models have been proposed to learn joint representations of multi-modal data, including natural language, visual signals, and vocal signals [298, 299]. Accordingly, an interesting future direction is how to effectively visualize learned joint representations of multi-modal data in an all-in-one manner, to facilitate the understanding of the data and their relationships. Various classic multi-modal tasks can 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.

6.3.2 Analyzing concept drift

In real-world applications, it is often assumed that the mapping from input data to output values (e.g., prediction label) is static. However, as data continues to arrive, the mapping between the input data and output values may change in unexpected ways [300]. In such a situation, a model trained on historical data may no longer work properly on new data. This usually causes noticeable performance degradation when the application data does not match the training data. Such a non-stationary learning problem over time is known as concept drift. As more and more machine learning applications directly consume streaming data, it is important to detect and analyze concept drift and minimize the resulting performance degradation [301, 302]. In the field of machine learning, three main research topics, have been studied: drift detection, drift understanding, and drift adaptation. Machine learning researchers have proposed many automatic algorithms to detect and adapt to 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 degree of drift at a given time. This makes it hard to understand why and where drift occurs. If the adaptation algorithms fail to improve the model performance, the black-box behavior of the adaptation models makes it difficult to diagnose the root cause of performance degradation. As a result, model developers need 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 to overcoming such drift. This requirement naturally leads to a visual analytics paradigm where the expert interacts and collaborates in concept drift detection and adaptation algorithms by putting the human in the loop. The major challenges here are how to (i) visually represent the evolving patterns of streaming data over time and effectively compare data distributions at different points in time, and (ii) tightly integrate such streaming data visualization with drift detection and adaptation algorithms to form an interactive and progressive analysis environment with the human in the loop.

7 Conclusions

This paper has comprehensively reviewed recent progress and developments in visual analytics techniques for machine learning. These techniques are classified into three groups by the corresponding analysis stage: techniques before, during, and after model building. Each category is detailed by typical analysis tasks, and each task is illustrated by a set of representative works. By comprehensively analyzing existing visual analytics research for machine learning, we also suggest six directions for 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 during model building, and multi-modal data understanding and concept drift analysis after model building. We hope this survey has provided an overview of visual analytics research for machine learning, facilitating understanding of state-of-the-art knowledge in this area, and shedding light on future research.

Acknowledgements

This research is supported by the National Key R&D Program of China (Nos. 2018YFB1004300 and 2019YFB1405703), the National Natural Science Foundation of China (Nos. 61761136020, 61672307, 61672308, and 61936002), TC190A4DA/3, the Institute Guo Qiang, Tsinghua University, and in part by Tsinghua–Kuaishou Institute of Future Media Data.

References

[1] Liu, S. X.; Wang, X. T.; Liu, M. C.; Zhu, J. Towards better analysis of machine learning models: A visual analytics perspective. *Visual Informatics* Vol. 1, No. 1, 48–56, 2017.

[2] Choo, J.; Liu, S. X. Visual analytics for explainable deep learning. *IEEE Computer Graphics and Applications* Vol. 38, No. 4, 84–92, 2018.

[3] Hohman, F.; Kahng, M.; Pienta, R.; Chau, D. H. Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 8, 2674–2693, 2019.

[4] Zeiler, M. D.; Fergus, R. Visualizing and understanding convolutional networks. In: *Computer Vision–ECCV 2014. Lecture Notes in Computer Science, Vol. 8689*. Fleet, D.; Pajdla, T.; Schiele, B.; Tuytelaars, T. Eds. Springer Cham, 818–833, 2014.

[5] Liu, S. X.; Wang, X. T.; Collins, C.; Dou, W. W.; Ouyang, F.; El-Assady, M.; Jiang, L.; Keim, D. A. Bridging text visualization and mining: A task-driven survey. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 7, 2482–2504, 2019.

[6] Lu, Y. F.; Garcia, R.; Hansen, B.; Gleicher, M.; Maciejewski, R. The state-of-the-art in predictive visual analytics. *Computer Graphics Forum* Vol. 36, No. 3, 539–562, 2017.

[7] Sacha, D.; Kraus, M.; Keim, D. A.; Chen, M. VIS4ML: An ontology for visual analytics assisted machine learning. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 385–395, 2019.

[8] Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; Batra, D. Grad-CAM: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision* Vol. 128, 336–359, 2020.

[9] Zhang, Q. S.; Zhu, S. C. Visual interpretability for deep learning: A survey. *Frontiers of Information Technology & Electronic Engineering* Vol. 19, No. 1, 27–39, 2018.

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

[11] Marsland, S. *Machine Learning: an Algorithmic Perspective*. Chapman and Hall/CRC, 2015.

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

[13] Choo, J.; Lee, C.; Reddy, C. K.; Park, H. UTOPIAN: User-driven topic modeling based on interactive nonnegative matrix factorization. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 1992–2001, 2013.

[14] Alemzadeh, S.; Niemann, U.; Ittermann, T.; Völzke, H.; Schneider, D.; Spiliopoulou, M.; Bühler, K.; Preim, B. Visual analysis of missing values in longitudinal cohort study data. *Computer Graphics Forum* Vol. 39, No. 1, 63–75, 2020.

[15] Arbesser, C.; Spechtenhauser, F.; Muhlbacher, T.; Piringer, H. Visplause: Visual data quality assessment of many time series using plausibility checks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 641–650, 2017.

[16] Bäuerle, A.; Neumann, H.; Ropinski, T. Classifier-guided visual correction of noisy labels for image classification tasks. *Computer Graphics Forum* Vol. 39, No. 3, 195–205, 2020.

[17] Bernard, J.; Hutter, M.; Reinemuth, H.; Pfeifer, H.; Bors, C.; Kohlhammer, J. Visual-interactive preprocessing of multivariate time series data. *Computer Graphics Forum* Vol. 38, No. 3, 401–412, 2019.

[18] Bernard, J.; Hutter, M.; Zeppelzauer, M.; Fellner, D.; Sedlmair, M. Comparing visual-interactive labeling with active learning: An experimental study. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 298–308, 2018.

[19] Bernard, J.; Zeppelzauer, M.; Lehmann, M.; Müller, M.; Sedlmair, M. Towards user-centered active learning algorithms. *Computer Graphics Forum* Vol. 37, No. 3, 121–132, 2018.

[20] Bors, C.; Gschwandtner, T.; Miksch, S. Capturing and visualizing provenance from data wrangling. *IEEE Computer Graphics and Applications* Vol. 39, No. 6, 61–75, 2019.

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

[22] Dextras-Romagnino, K.; Munzner, T. Segmen++ tifier: Interactive refinement of clickstream data. *Computer Graphics Forum* Vol. 38, No. 3, 623–634, 2019.

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

[24] Halter, G.; Ballester-Ripoll, R.; Flueckiger, B.; Pajarola, R. VIAN: A visual annotation tool for film analysis. *Computer Graphics Forum* Vol. 38, No. 3, 119–129, 2019.

[25] Heimerl, F.; Koch, S.; Bosch, H.; Ertl, T. Visual classifier training for text document retrieval. *IEEE Transactions on Visualization and Computer Graphics* Vol. 18, No. 12, 2839–2848, 2012.

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

[27] Soares Junior, A.; Renso, C.; Matwin, S. ANALYTiC: An active learning system for trajectory classification. *IEEE Computer Graphics and Applications* Vol. 37, No. 5, 28–39, 2017.

[28] Khayat, M.; Karimzadeh, M.; Zhao, J. Q.; Ebert, D. S. VASSL: A visual analytics toolkit for social spambot labeling. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 874–883, 2020.

[29] Kurzhals, K.; Hlawatsch, M.; Seeger, C.; Weiskopf, D. Visual analytics for mobile eye tracking. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 301–310, 2017.

[30] Lekschas, F.; Peterson, B.; Haehn, D.; Ma, E.; Gehlenborg, N.; Pfister, H. 2019. PEAX: interactive visual pattern search in sequential data using unsupervised deep representation learning. bioRxiv 597518, https://doi.org/10.1101/597518, 2020.

[31] Liu, S. X.; Chen, C. J.; Lu, Y. F.; Ouyang, F. X.; Wang, B. An interactive method to improve crowdsourced annotations. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 235–245, 2019.

[32] Moehrmann, J.; Bernstein, S.; Schlegel, T.; Werner, G.; Heidemann, G. Improving the usability of hierarchical representations for interactively labeling large image data sets. In: *Human-Computer Interaction. Design and Development Approaches. Lecture Notes in Computer Science, Vol. 6761.* Jacko, J. A. Ed. Springer Berlin, 618–627, 2011.

[33] Paiva, J. G. S.; Schwartz, W. R.; Pedrini, H.; Minghim, R. An approach to supporting incremental visual data classification. *IEEE Transactions on*

Visualization and Computer Graphics Vol. 21, No. 1, 4–17, 2015.

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

[35] Park, J. H.; Nadeem, S.; Mirhosseini, S.; Kaufman, A. C2A: Crowd consensus analytics for virtual colonoscopy. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 21–30, 2016.

[36] De Rooij, O.; van Wijk, J. J.; Worring, M. MediaTable: Interactive categorization of multimedia collections. *IEEE Computer Graphics and Applications* Vol. 30, No. 5, 42–51, 2010.

[37] Snyder, L. S.; Lin, Y. S.; Karimzadeh, M.; Goldwasser, D.; Ebert, D. S. Interactive learning for identifying relevant tweets to support real-time situational awareness. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 558–568, 2020.

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

[39] Stein, M.; Janetzko, H.; Breitkreutz, T.; Seebacher, D.; Schreck, T.; Grossniklaus, M.; Couzin, I. D.; Keim, D. A. Director's cut: Analysis and annotation of soccer matches. *IEEE Computer Graphics and Applications* Vol. 36, No. 5, 50–60, 2016.

[40] Wang, X. M.; Chen, W.; Chou, J. K.; Bryan, C.; Guan, H. H.; Chen, W. L.; Pan, R.; Ma, K.-L. GraphProtector: A visual interface for employing and assessing multiple privacy preserving graph algorithms. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 193–203, 2019.

[41] Wang, X. M.; Chou, J. K.; Chen, W.; Guan, H. H.; Chen, W. L.; Lao, T. Y.; Ma, K.-L. A utility-aware visual approach for anonymizing multi-attribute tabular data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 351–360, 2018.

[42] Willett, W.; Ginosar, S.; Steinitz, A.; Hartmann, B.; Agrawala, M. Identifying redundancy and exposing provenance in crowdsourced data analysis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2198–2206, 2013.

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

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

[45] Krause, J.; Perer, A.; Bertini, E. INFUSE: Interactive feature selection for predictive modeling of high dimensional data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1614–1623, 2014.

[46] May, T.; Bannach, A.; Davey, J.; Ruppert, T.; Kohlhammer, J. Guiding feature subset selection with an interactive visualization. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 111–120, 2011.

[47] Muhlbacher, T.; Piringer, H. A partition-based framework for building and validating regression models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 1962–1971, 2013.

[48] Seo, J.; Shneiderman, B. A rank-by-feature framework for interactive exploration of multidimensional data. *Information Visualization* Vol. 4, No. 2, 96–113, 2005.

[49] Tam, G. K. L.; Fang, H.; Aubrey, A. J.; Grant, P. W.; Rosin, P. L.; Marshall, D.; Chen, M. Visualization of time-series data in parameter space for understanding facial dynamics. *Computer Graphics Forum* Vol. 30, No. 3, 901–910, 2011.

[50] Broeksema, B.; Baudel, T.; Telea, A.; Crisafulli, P. Decision exploration lab: A visual analytics solution for decision management. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 1972–1981, 2013.

[51] Cashman, D.; Patterson, G.; Mosca, A.; Watts, N.; Robinson, S.; Chang, R. RNNbow: Visualizing learning via backpropagation gradients in RNNs. *IEEE Computer Graphics and Applications* Vol. 38, No. 6, 39–50, 2018.

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

[53] Eichner, C.; Schumann, H.; Tominski, C. Making parameter dependencies of time-series segmentation visually understandable. *Computer Graphics Forum* Vol. 39, No. 1, 607–622, 2020.

[54] Ferreira, N.; Lins, L.; Fink, D.; Kelling, S.; Wood, C.; Freire, J.; Silva, C. BirdVis: Visualizing and

understanding bird populations. *IEEE Transactions on Visualization and Computer Graphics* Vol. 17, No. 12, 2374–2383, 2011.

[55] Fröhler, B.; Möller, T.; Heinzl, C. GEMSe: Visualization-guided exploration of multi-channel segmentation algorithms. *Computer Graphics Forum* Vol. 35, No. 3, 191–200, 2016.

[56] Hohman, F.; Park, H.; Robinson, C.; Polo Chau, D. H. Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 1096–1106, 2020.

[57] Jaunet, T.; Vuillemot, R.; Wolf, C. DRLViz: Understanding decisions and memory in deep reinforcement learning. *Computer Graphics Forum* Vol. 39, No. 3, 49–61, 2020.

[58] Jean, C. S.; Ware, C.; Gamble, R. Dynamic change arcs to explore model forecasts. *Computer Graphics Forum* Vol. 35, No. 3, 311–320, 2016.

[59] Kahng, M.; Andrews, P. Y.; Kalro, A.; Chau, D. H. ActiVis: Visual exploration of industry-scale deep neural network models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 88–97, 2018.

[60] Kahng, M.; Thorat, N.; Chau, D. H. P.; Viegas, F. B.; Wattenberg, M. GAN lab: Understanding complex deep generative models using interactive visual experimentation. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 310–320, 2019.

[61] Kwon, B. C.; Anand, V.; Severson, K. A.; Ghosh, S.; Sun, Z. N.; Frohnert, B. I.; Lundgren, M.; Ng, K. DPVis: Visual analytics with hidden Markov models for disease progression pathways. *IEEE Transactions on Visualization and Computer Graph*ics doi: 10.1109/TVCG.2020.2985689, 2020.

[62] Liu, M. C.; Shi, J. X.; Li, Z.; Li, C. X.; Zhu, J.; Liu, S. X. Towards better analysis of deep convolutional neural networks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 91–100, 2017.

[63] Liu, S. S.; Li, Z. M.; Li, T.; Srikumar, V.; Pascucci, V.; Bremer, P. T. NLIZE: A perturbation-driven visual interrogation tool for analyzing and interpreting natural language inference models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 651–660, 2019.

[64] Migut, M.; van Gemert, J.; Worring, M. Interactive decision making using dissimilarity to visually represented prototypes. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 141–149, 2011.

[65] Ming, Y.; Cao, S.; Zhang, R.; Li, Z.; Chen, Y.; Song, Y.; Qu, H. Understanding hidden memories of recurrent neural networks. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 13–24, 2017.

[66] Ming, Y.; Qu, H. M.; Bertini, E. RuleMatrix: Visualizing and understanding classifiers with rules. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 342–352, 2019.

[67] Murugesan, S.; Malik, S.; Du, F.; Koh, E.; Lai, T. M. DeepCompare: Visual and interactive comparison of deep learning model performance. *IEEE Computer Graphics and Applications* Vol. 39, No. 5, 47–59, 2019.

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

[69] Rauber, P. E.; Fadel, S. G.; Falcao, A. X.; Telea, A. C. Visualizing the hidden activity of artificial neural networks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 101–110, 2017.

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

[71] Scheepens, R.; Michels, S.; van de Wetering, H.; van Wijk, J. J. Rationale visualization for safety and security. *Computer Graphics Forum* Vol. 34, No. 3, 191–200, 2015.

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

[73] Strobelt, H.; Gehrmann, S.; Pfister, H.; Rush, A. M. LSTMVis: A tool for visual analysis of hidden state dynamics in recurrent neural networks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 667–676, 2018.

[74] Wang, J. P.; Gou, L.; Yang, H.; Shen, H. W. GANViz: A visual analytics approach to understand the adversarial game. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 6, 1905–1917, 2018.

[75] Wang, J. P.; Gou, L.; Zhang, W.; Yang, H.; Shen, H. W. DeepVID: Deep visual interpretation and diagnosis for image classifiers via knowledge distillation. *IEEE*

Transactions on Visualization and Computer Graphics Vol. 25, No. 6, 2168–2180, 2019.

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

[77] Wongsuphasawat, K.; Smilkov, D.; Wexler, J.; Wilson, J.; Mane, D.; Fritz, D.; Krishnan, D.; Viegas, F. B.; Wattenberg, M. Visualizing dataflow graphs of deep learning models in TensorFlow. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 1–12, 2018.

[78] Zhang, C.; Yang, J.; Zhan, F. B.; Gong, X.; Brender, J. D.; Langlois, P. H.; Barlowe, S.; Zhao, Y. A visual analytics approach to high-dimensional logistic regression modeling and its application to an environmental health study. In: Proceedings of the IEEE Pacific Visualization Symposium, 136–143, 2016.

[79] Zhao, X.; Wu, Y. H.; Lee, D. L.; Cui, W. W. iForest: Interpreting random forests via visual analytics. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 407–416, 2019.

[80] Ahn, Y.; Lin, Y. R. FairSight: Visual analytics for fairness in decision making. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 1086–1095, 2019.

[81] Alsallakh, B.; Hanbury, A.; Hauser, H.; Miksch, S.; Rauber, A. Visual methods for analyzing probabilistic classification data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1703–1712, 2014.

[82] Bilal, A.; Jourabloo, A.; Ye, M.; Liu, X. M.; Ren, L. 2018. Do convolutional neural networks learn class hierarchy? *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 152–162, 2018.

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

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

[85] Diehl, A.; Pelorosso, L.; Delrieux, C.; Matković, K.; Ruiz, J.; Gröller, M. E.; Bruckner, S. Albero: A visual analytics approach for probabilistic weather forecasting. *Computer Graphics Forum* Vol. 36, No. 7, 135–144, 2017.

[86] Gleicher, M.; Barve, A.; Yu, X. Y.; Heimerl, F. Boxer: Interactive comparison of classifier results. *Computer Graphics Forum* Vol. 39, No. 3, 181–193, 2020.

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

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

[89] Liu, M. C.; Shi, J. X.; Cao, K. L.; Zhu, J.; Liu, S. X. Analyzing the training processes of deep generative models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 77–87, 2018.

[90] Liu, S. X.; Xiao, J. N.; Liu, J. L.; Wang, X. T.; Wu, J.; Zhu, J. Visual diagnosis of tree boosting methods. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 163–173, 2018.

[91] Ma, Y. X.; Xie, T. K.; Li, J. D.; Maciejewski, R. Explaining vulnerabilities to adversarial machine learning through visual analytics. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 1075–1085, 2020.

[92] Pezzotti, N.; Hollt, T.; van Gemert, J.; Lelieveldt, B. P. F.; Eisemann, E.; Vilanova, A. DeepEyes: Progressive visual analytics for designing deep neural networks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 98–108, 2018.

[93] Ren, D. H.; Amershi, S.; Lee, B.; Suh, J.; Williams, J. D. Squares: Supporting interactive performance analysis for multiclass classifiers. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 61–70, 2017.

[94] Spinner, T.; Schlegel, U.; Schafer, H.; El-Assady, M. explAIner: A visual analytics framework for interactive and explainable machine learning. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 1064–1074, 2020.

[95] Strobelt, H.; Gehrmann, S.; Behrisch, M.; Perer, A.; Pfister, H.; Rush, A. M. Seq2seq-Vis: A visual debugging tool for sequence-to-sequence models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 353–363, 2019.

[96] Wang, J. P.; Gou, L.; Shen, H. W.; Yang, H. DQNViz: A visual analytics approach to understand deep Q-networks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 288–298, 2019.

[97] Wexler, J.; Pushkarna, M.; Bolukbasi, T.; Wattenberg, M.; Viegas, F.; Wilson, J. The what-if tool: Interactive probing of machine learning models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 56–65, 2019.

[98] Zhang, J. W.; Wang, Y.; Molino, P.; Li, L. Z.; Ebert, D. S. Manifold: A model-agnostic framework for interpretation and diagnosis of machine learning models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 364–373, 2019.

[99] Bogl, M.; Aigner, W.; Filzmoser, P.; Lammarsch, T.; Miksch, S.; Rind, A. Visual analytics for model selection in time series analysis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2237–2246, 2013.

[100] Cashman, D.; Perer, A.; Chang, R.; Strobelt, H. Ablate, variate, and contemplate: Visual analytics for discovering neural architectures. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 863–873, 2020.

[101] Cavallo, M.; Demiralp, Ç. Track xplorer: A system for visual analysis of sensor-based motor activity predictions. *Computer Graphics Forum* Vol. 37, No. 3, 339–349, 2018.

[102] Cavallo, M.; Demiralp, C. Clustrophile 2: Guided visual clustering analysis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 267–276, 2019.

[103] Das, S.; Cashman, D.; Chang, R.; Endert, A. BEAMES: Interactive multimodel steering, selection, and inspection for regression tasks. *IEEE Computer Graphics and Applications* Vol. 39, No. 5, 20–32, 2019.

[104] Dingen, D.; van't Veer, M.; Houthuizen, P.; Mestrom, E. H. J.; Korsten, E. H. H. M.; Bouwman, A. R. A.; van Wijk. J. J. RegressionExplorer: Interactive exploration of logistic regression models with subgroup analysis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 246–255, 2019.

[105] Dou, W. W.; Yu, L.; Wang, X. Y.; Ma, Z. Q.; Ribarsky, W. HierarchicalTopics: Visually exploring large text collections using topic hierarchies. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2002–2011, 2013.

[106] El-Assady, M.; Kehlbeck, R.; Collins, C.; Keim, D.; Deussen, O. Semantic concept spaces: Guided topic model refinement using word-embedding projections. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 1001–1011, 2020.

[107] El-Assady, M.; Sevastjanova, R.; Sperrle, F.; Keim, D.; Collins, C. Progressive learning of topic modeling parameters: A visual analytics framework. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 382–391, 2018.

[108] El-Assady, M.; Sperrle, F.; Deussen, O.; Keim, D.; Collins, C. Visual analytics for topic model optimization based on user-steerable speculative execution. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 374–384, 2019.

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

[110] Kwon, B. C.; Choi, M. J.; Kim, J. T.; Choi, E.; Kim, Y. B.; Kwon, S.; Sun, J.; Choo, J. RetainVis: Visual analytics with interpretable and interactive recurrent neural networks on electronic medical records. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 299–309, 2019.

[111] Lee, H.; Kihm, J.; Choo, J.; Stasko, J.; Park, H. iVisClustering: An interactive visual document clustering via topic modeling. *Computer Graphics Forum* Vol. 31, No. 3, 1155–1164, 2012.

[112] Liu, M. C.; Liu, S. X.; Zhu, X. Z.; Liao, Q. Y.; Wei, F. R.; Pan, S. M. An uncertainty-aware approach for exploratory microblog retrieval. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 250–259, 2016.

[113] Lowe, T.; Forster, E. C.; Albuquerque, G.; Kreiss, J. P.; Magnor, M. Visual analytics for development and evaluation of order selection criteria for autoregressive processes. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 151–159, 2016.

[114] MacInnes, J.; Santosa, S.; Wright, W. Visual classification: Expert knowledge guides machine learning. *IEEE Computer Graphics and Applications* Vol. 30, No. 1, 8–14, 2010.

[115] Migut, M.; Worring, M. Visual exploration of classification models for risk assessment. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 11–18, 2010.

[116] Ming, Y.; Xu, P. P.; Cheng, F. R.; Qu, H. M.; Ren, L. ProtoSteer: Steering deep sequence model with prototypes. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 238–248, 2020.

[117] Muhlbacher, T.; Linhardt, L.; Moller, T.; Piringer, H. TreePOD: Sensitivity-aware selection of Pareto-optimal decision trees. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 174–183, 2018.

[118] Packer, E.; Bak, P.; Nikkila, M.; Polishchuk, V.; Ship, H. J. Visual analytics for spatial clustering: Using a heuristic approach for guided exploration. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2179–2188, 2013.

[119] Piringer, H.; Berger, W.; Krasser, J. HyperMoVal: Interactive visual validation of regression models for real-time simulation. *Computer Graphics Forum* Vol. 29, No. 3, 983–992, 2010.

[120] Sacha, D.; Kraus, M.; Bernard, J.; Behrisch, M.; Schreck, T.; Asano, Y.; Keim, D. A. SOMFlow: Guided exploratory cluster analysis with self-organizing maps and analytic provenance. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 120–130, 2018.

[121] Schultz, T.; Kindlmann, G. L. Open-box spectral clustering: Applications to medical image analysis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2100–2108, 2013.

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

[123] Vrotsou, K.; Nordman, A. Exploratory visual sequence mining based on pattern-growth. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 8, 2597–2610, 2019.

[124] Wang, X. T.; Liu, S. X.; Liu, J. L.; Chen, J. F.; Zhu, J.; Guo, B. N. TopicPanorama: A full picture of relevant topics. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 12, 2508–2521, 2016.

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

[126] Zhao, K. Y.; Ward, M. O.; Rundensteiner, E. A.; Higgins, H. N. LoVis: Local pattern visualization for model refinement. *Computer Graphics Forum* Vol. 33, No. 3, 331–340, 2014.

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

[128] Berger, M.; McDonough, K.; Seversky, L. M. Cite2vec: Citation-driven document exploration via word embeddings. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 691–700, 2017.

[129] Blumenschein, M.; Behrisch, M.; Schmid, S.; Butscher, S.; Wahl, D. R.; Villinger, K.; Renner, B.; Reiterer, H.; Keim, D. A. 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, 36–47, 2018.

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

[131] Broeksema, B.; Telea, A. C.; Baudel, T. Visual analysis of multi-dimensional categorical data sets. *Computer Graphics Forum* Vol. 32, No. 8, 158–169, 2013.

[132] Cao, N.; Sun, J. M.; Lin, Y. R.; Gotz, D.; Liu, S. X.; Qu, H. M. FacetAtlas: Multifaceted visualization for rich text corpora. *IEEE Transactions on Visualization and Computer Graphics* Vol. 16, No. 6, 1172–1181, 2010.

[133] Chandrasegaran, S.; Badam, S. K.; Kisselburgh, L.; Ramani, K.; Elmqvist, N. Integrating visual analytics support for grounded theory practice in qualitative text analysis. *Computer Graphics Forum* Vol. 36, No. 3, 201–212, 2017.

[134] Chen, S. M.; Andrienko, N.; Andrienko, G.; Adilova, L.; Barlet, J.; Kindermann, J.; Nguyen, P. H.; Thonnard, O.; Turkay, C. LDA ensembles for interactive exploration and categorization of behaviors. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 9, 2775–2792, 2020.

[135] Correll, M.; Witmore, M.; Gleicher, M. Exploring collections of tagged text for literary scholarship. *Computer Graphics Forum* Vol. 30, No. 3, 731–740, 2011.

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

[137] El-Assady, M.; Gold, V.; Acevedo, C.; Collins, C.; Keim, D. ConToVi: Multi-party conversation exploration using topic-space views. *Computer Graphics Forum* Vol. 35, No. 3, 431–440, 2016.

[138] El-Assady, M.; Sevastjanova, R.; Keim, D.; Collins, C. ThreadReconstructor: Modeling reply-chains to untangle conversational text through visual analytics. *Computer Graphics Forum* Vol. 37, No. 3, 351–365, 2018.

[139] Filipov, V.; Arleo, A.; Federico, P.; Miksch, S. CV3: Visual exploration, assessment, and comparison of CVs. *Computer Graphics Forum* Vol. 38, No. 3, 107–118, 2019.

[140] Fried, D.; Kobourov, S. G. Maps of computer science. In: Proceedings of the IEEE Pacific Visualization Symposium, 113–120, 2014.

[141] Fulda, J.; Brehmer, M.; Munzner, T. TimeLineCurator: Interactive authoring of visual timelines from unstructured text. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 300–309, 2016.

[142] Glueck, M.; Naeini, M. P.; Doshi-Velez, F.; Chevalier, F.; Khan, A.; Wigdor, D.; Brudno, M. PhenoLines: Phenotype comparison visualizations for disease subtyping via topic models. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 371–381, 2018.

[143] Gorg, C.; Liu, Z. C.; Kihm, J.; Choo, J.; Park, H.; Stasko, J. Combining computational analyses and interactive visualization for document exploration and sensemaking in jigsaw. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 10, 1646–1663, 2013.

[144] Guo, H.; Laidlaw, D. H. Topic-based exploration and embedded visualizations for research idea generation. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 3, 1592–1607, 2020.

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

[146] Hong, F.; Lai, C.; Guo, H.; Shen, E.; Yuan, X.; Li. S. FLDA: Latent Dirichlet allocation based unsteady flow analysis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No.12, 2545–2554, 2014.

[147] Hoque, E.; Carenini, G. ConVis: A visual text analytic system for exploring blog conversations. *Computer Graphics Forum* Vol. 33, No. 3, 221–230, 2014.

[148] Hu, M. D.; Wongsuphasawat, K.; Stasko, J. Visualizing social media content with SentenTree. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 621–630, 2017.

[149] Jänicke, H.; Borgo, R.; Mason, J. S. D.; Chen, M. SoundRiver: Semantically-rich sound illustration. *Computer Graphics Forum* Vol. 29, No. 2, 357–366, 2010.

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

[151] Jankowska, M.; Kefiselj, V.; Milios, E. Relative N-gram signatures: Document visualization at the level of character n-grams. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 103–112, 2012.

[152] Ji, X. N.; Shen, H. W.; Ritter, A.; Machiraju, R.; Yen, P. Y. Visual exploration of neural document embedding in information retrieval: Semantics and feature selection. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 6, 2181–2192, 2019.

[153] Kakar, T.; Qin, X.; Rundensteiner, E. A.; Harrison, L.; Sahoo, S. K.; De, S. DIVA: Exploration and validation of hypothesized drug-drug interactions. *Computer Graphics Forum* Vol. 38, No. 3, 95–106, 2019.

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

[155] Kim, M.; Kang, K.; Park, D.; Choo, J.; Elmqvist, N. TopicLens: Efficient multi-level visual topic exploration of large-scale document collections. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 151–160, 2017.

[156] Kochtchi, A.; von Landesberger, T.; Biemann, C. Networks of names: Visual exploration and semi-automatic tagging of social networks from newspaper articles. *Computer Graphics Forum* Vol. 33, No. 3, 211–220, 2014.

[157] Li, M. Z.; Choudhury, F.; Bao, Z. F.; Samet, H.; Sellis, T. ConcaveCubes: Supporting cluster-based geographical visualization in large data scale. *Computer Graphics Forum* Vol. 37, No. 3, 217–228, 2018.

[158] Liu, S.; Wang, B.; Thiagarajan, J. J.; Bremer, P. T.; Pascucci, V. Visual exploration of high-dimensional data through subspace analysis and dynamic projections. *Computer Graphics Forum* Vol. 34, No. 3, 271–280, 2015.

[159] Liu, S.; Wang, X.; Chen, J.; Zhu, J.; Guo, B. TopicPanorama: A full picture of relevant topics. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 183–192, 2014.

[160] Liu, X.; Xu, A.; Gou, L.; Liu, H.; Akkiraju, R.; Shen, H. W. SocialBrands: Visual analysis of public perceptions of brands on social media. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 71–80, 2016.

[161] Oelke, D.; Strobelt, H.; Rohrdantz, C.; Gurevych, I.; Deussen, O. Comparative exploration of document collections: A visual analytics approach. *Computer Graphics Forum* Vol. 33, No. 3, 201–210, 2014.

[162] Park, D.; Kim, S.; Lee, J.; Choo, J.; Diakopoulos, N.; Elmqvist, N. ConceptVector: text visual analytics via interactive lexicon building using word embedding. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 361–370, 2018.

[163] Paulovich, F. V.; Toledo, F. M. B.; Telles, G. P.; Minghim, R.; Nonato, L. G. Semantic wordification of document collections. *Computer Graphics Forum* Vol. 31, No. 3pt3, 1145–1153, 2012.

[164] Shen, Q. M.; Zeng, W.; Ye, Y.; Arisona, S. M.; Schubiger, S.; Burkhard, R.; Qu, H. StreetVizor: Visual exploration of human-scale urban forms based on street views. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 1004–1013, 2018.

[165] Von Landesberger, T.; Basgier, D.; Becker, M. Comparative local quality assessment of 3D medical image segmentations with focus on statistical shape model-based algorithms. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 12, 2537–2549, 2016.

[166] Wall, E.; Das, S.; Chawla, R.; Kalidindi, B.; Brown, E. T.; Endert, A. Podium: Ranking data using mixed-initiative visual analytics. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 288–297, 2018.

[167] Xie, X.; Cai, X. W.; Zhou, J. P.; Cao, N.; Wu, Y. C. A semantic-based method for visualizing large image collections. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 7, 2362–2377, 2019.

[168] Zhang, L.; Huang, H. Hierarchical narrative collage for digital photo album. *Computer Graphics Forum* Vol. 31, No. 7, 2173–2181, 2012.

[169] Zhao, J.; Chevalier, F.; Collins, C.; Balakrishnan, R. Facilitating discourse analysis with interactive visualization. *IEEE Transactions on Visualization and Computer Graphics* Vol. 18, No. 12, 2639–2648, 2012.

[170] Alsakran, J.; Chen, Y.; Luo, D. N.; Zhao, Y.; Yang, J.; Dou, W. W.; Liu, S. Real-time visualization of streaming text with a force-based dynamic system. *IEEE Computer Graphics and Applications* Vol. 32, No. 1, 34–45, 2012.

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

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

[173] Andrienko, G.; Andrienko, N.; Bremm, S.; Schreck, T.; von Landesberger, T.; Bak, P.; Keim, D. Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns. *Computer Graphics Forum* Vol. 29, No. 3, 913–922, 2010.

[174] Andrienko, G.; Andrienko, N.; Hurter, C.; Rinzivillo, S.; Wrobel, S. Scalable analysis of movement data for extracting and exploring significant places. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 7, 1078–1094, 2013.

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

[176] Bögl, M.; Filzmoser, P.; Gschwandtner, T.; Lammarsch, T.; Leite, R. A.; Miksch, S.; Rind, A. Cycle plot revisited: Multivariate outlier detection using a distance-based abstraction. *Computer Graphics Forum* Vol. 36, No. 3, 227–238, 2017.

[177] Bosch, H.; Thom, D.; Heimerl, F.; Puttmann, E.; Koch, S.; Kruger, R.; Worner, M.; Ertl, T. ScatterBlogs2: real-time monitoring of microblog messages through user-guided filtering. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2022–2031, 2013.

[178] Buchmüller, J.; Janetzko, H.; Andrienko, G.; Andrienko, N.; Fuchs, G.; Keim, D. A. Visual analytics for exploring local impact of air traffic. *Computer Graphics Forum* Vol. 34, No. 3, 181–190, 2015.

[179] Cao, N.; Lin, C. G.; Zhu, Q. H.; Lin, Y. R.; Teng, X.; Wen, X. D. Voila: Visual anomaly detection and monitoring with streaming spatiotemporal data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 23–33, 2018.

[180] Cao, N.; Lin, Y. R.; Sun, X. H.; Lazer, D.; Liu, S. X.; Qu, H. M. Whisper: Tracing the spatiotemporal process of information diffusion in real time. *IEEE Transactions on Visualization and Computer Graphics* Vol. 18, No. 12, 2649–2658, 2012.

[181] Cao, N.; Shi, C. L.; Lin, S.; Lu, J.; Lin, Y. R.; Lin, C. Y. TargetVue: Visual analysis of anomalous user behaviors in online communication systems. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 280–289, 2016.

[182] Chae, J.; Thom, D.; Bosch, H.; Jang, Y.; Maciejewski, R.; Ebert, D. S.; Ertl, T. 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, 143–152, 2012.

[183] Chen, Q.; Yue, X. W.; Plantaz, X.; Chen, Y. Z.; Shi, C. L.; Pong, T. C.; Qu, H. ViSeq: Visual analytics of learning sequence in massive open online courses. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 3, 1622–1636, 2020.

[184] Chen, S.; Chen, S.; Lin, L.; Yuan, X.; Liang, J.; Zhang, X. 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, 36–47, 2017.

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

[186] Chen, S. M.; Yuan, X. R.; Wang, Z. H.; Guo, C.; Liang, J.; Wang, Z. C.; Zhang, X.; Zhang, J. Interactive visual discovering of movement patterns from sparsely sampled geo-tagged social media data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 270–279, 2016.

[187] Chen, Y.; Chen, Q.; Zhao, M.; Boyer, S.; Veeramachaneni, K.; Qu, H. DropoutSeer: Visualizing learning patterns in massive open online courses for dropout reasoning and prediction. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 111–120, 2016.

[188] Chen, Y. Z.; Xu, P. P.; Ren, L. Sequence synopsis: Optimize visual summary of temporal event data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 45–55, 2018.

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

[190] Cui, W. W.; Liu, S. X.; Tan, L.; Shi, C. L.; Song, Y. Q.; Gao, Z. K.; Qu, H. M.; Tong, X. TextFlow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics* Vol. 17, No. 12, 2412–2421, 2011.

[191] Cui, W. W.; Liu, S. X.; Wu, Z. F.; Wei, H. How hierarchical topics evolve in large text corpora. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 2281–2290, 2014.

[192] Di Lorenzo, G.; Sbodio, M.; Calabrese, F.; Berlingerio, M.; Pinelli, F.; Nair, R. AllAboard: Visual exploration of cellphone mobility data to optimise public transport. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 2, 1036–1050, 2016.

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

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

[195] Du, F.; Plaisant, C.; Spring, N.; Shneiderman, B. EventAction: Visual analytics for temporal event sequence recommendation. In: Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 61–70, 2016.

[196] El-Assady, M.; Sevastjanova, R.; Gipp, B.; Keim, D.; Collins, C. NEREx: Named-entity relationship exploration in multi-party conversations. *Computer Graphics Forum* Vol. 36, No. 3, 213–225, 2017.

[197] Fan, M. M.; Wu, K.; Zhao, J.; Li, Y.; Wei, W.; Truong, K. N. VisTA: Integrating machine intelligence with visualization to support the investigation of think-aloud sessions. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 343–352, 2020.

[198] Ferreira, N.; Poco, J.; Vo, H. T.; Freire, J.; Silva, C. T. Visual exploration of big spatio-temporal urban data: A study of New York City taxi trips. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2149–2158, 2013.

[199] Gobbo, B.; Balsamo, D.; Mauri, M.; Bajardi, P.; Panisson, A.; Ciuccarelli, P. Topic Tomographies (TopTom): A visual approach to distill information from media streams. *Computer Graphics Forum* Vol. 38, No. 3, 609–621, 2019.

[200] Gotz, D.; Stavropoulos, H. DecisionFlow: Visual analytics for high-dimensional temporal event sequence data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1783–1792, 2014.

[201] Guo, S. N.; Jin, Z. C.; Gotz, D.; Du, F.; Zha, H. Y.; Cao, N. Visual progression analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 417–426, 2019.

[202] Guo, S. N.; Xu, K.; Zhao, R. W.; Gotz, D.; Zha, H. Y.; Cao, N. EventThread: Visual summarization and stage analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 56–65, 2018.

[203] Gutenko, I.; Dmitriev, K.; Kaufman, A. E.; Barish, M. A. AnaFe: Visual analytics of image-derived temporal features: Focusing on the spleen. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 171–180, 2017.

[204] Havre, S.; Hetzler, E.; Whitney, P.; Nowell, L. ThemeRiver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics* Vol. 8, No. 1, 9–20, 2002.

[205] Heimerl, F.; Han, Q.; Koch, S.; Ertl, T. CiteRivers: Visual analytics of citation patterns. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 190–199, 2016.

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

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

[208] Kamaleswaran, R.; Collins, C.; James, A.; McGregor, C. PhysioEx: Visual analysis of physiological event streams. *Computer Graphics Forum* Vol. 35, No. 3, 331–340, 2016.

[209] Karduni, A.; Cho, I.; Wessel, G.; Ribarsky, W.; Sauda, E.; Dou, W. W. Urban space explorer: A visual analytics system for urban planning. *IEEE Computer Graphics and Applications* Vol. 37, No. 5, 50–60, 2017.

[210] Krueger, R.; Han, Q.; Ivanov, N.; Mahtal, S.; Thom, D.; Pfister, H.; Ertl, T. Bird's-eye-large-scale visual analytics of city dynamics using social location data. *Computer Graphics Forum* Vol. 38, No. 3, 595–607, 2019.

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

[212] Krueger, R.; Thom, D.; Ertl, T. Semantic enrichment of movement behavior with foursquare—A visual analytics approach. *IEEE Transactions on Visualization and Computer Graphics* Vol. 21, No. 8, 903–915, 2015.

[213] Lee, C.; Kim, Y.; Jin, S.; Kim, D.; Maciejewski, R.; Ebert, D.; Ko, S. A visual analytics system for exploring, monitoring, and forecasting road traffic congestion. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 11, 3133–3146, 2020.

[214] Leite, R. A.; Gschwandtner, T.; Miksch, S.; Kriglstein, S.; Pohl, M.; Gstrein, E.; Kuntner, J. EVA: Visual analytics to identify fraudulent events. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 1, 330–339, 2018.

[215] Li, J.; Chen, S. M.; Chen, W.; Andrienko, G.; Andrienko, N. Semantics-space-time cube. A conceptual framework for systematic analysis of texts in space and time. *IEEE Transactions on Visualization and Computer Graphics*, Vol. 26, No. 4, 1789–1806, 2019.

[216] Li, Q.; Wu, Z. M.; Yi, L. L.; Kristanto, S. N.; Qu, H. M.; Ma, X. J. WeSeer: Visual analysis for better information cascade prediction of WeChat articles. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 2, 1399–1412, 2020.

[217] Li, Z. Y.; Zhang, C. H.; Jia, S. C.; Zhang, J. W. Galex: Exploring the evolution and intersection of disciplines. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 1182–1192, 2019.

[218] Liu, H.; Jin, S. C.; Yan, Y. Y.; Tao, Y. B.; Lin, H. Visual analytics of taxi trajectory data via topical sub-trajectories. *Visual Informatics* Vol. 3, No. 3, 140–149, 2019.

[219] Liu, S. X.; Yin, J. L.; Wang, X. T.; Cui, W. W.; Cao, K. L.; Pei, J. Online visual analytics of text streams. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 11, 2451–2466, 2016.

[220] Liu, S.; Zhou, M. X.; Pan, S.; Song, Y.; Qian, W.; Cai, W.; Lian, X. TIARA: Interactive, topic-based visual text summarization and analysis. *ACM Transactions on Intelligent Systems and Technology* Vol. 3, No.2, Article No. 25, 2012.

[221] Liu, Z. C.; Kerr, B.; Dontcheva, M.; Grover, J.; Hoffman, M.; Wilson, A. CoreFlow: Extracting and visualizing branching patterns from event sequences. *Computer Graphics Forum* Vol. 36, No. 3, 527–538, 2017.

[222] Liu, Z.; Wang, Y.; Dontcheva, M.; Hofiman, M.; Walker, S.; Wilson, A. Patterns and sequences: Interactive exploration of clickstreams to understand common visitor paths. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No.1, 321–330, 2017.

[223] Lu, Y. F.; Steptoe, M.; Burke, S.; Wang, H.; Tsai, J. Y.; Davulcu, H.; Montgomery, D.; Corman, S. R.; Maciejewski, R. Exploring evolving media discourse through event cueing. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 220–229, 2016.

[224] Lu, Y. F.; Wang, F.; Maciejewski, R. Business intelligence from social media: A study from the VAST box office challenge. *IEEE Computer Graphics and Applications* Vol. 34, No. 5, 58–69, 2014.

[225] Lu, Y. F.; Wang, H.; Landis, S.; Maciejewski, R. A visual analytics framework for identifying topic drivers in media events. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 9, 2501–2515, 2018.

[226] Luo, D. N.; Yang, J.; Krstajic, M.; Ribarsky, W.; Keim, D. A. EventRiver: Visually exploring text collections with temporal references. *IEEE Transactions on Visualization and Computer Graphics* Vol. 18, No. 1, 93–105, 2012.

[227] Maciejewski, R.; Hafen, R.; Rudolph, S.; Larew, S. G.; Mitchell, M. A.; Cleveland, W. S.; Ebert, D. S. Forecasting hotspots: A predictive analytics approach. *IEEE Transactions on Visualization and Computer Graphics* Vol. 17, No. 4, 440–453, 2011.

[228] Malik, A.; Maciejewski, R.; Towers, S.; McCullough, S.; Ebert, D. S. Proactive spatiotemporal resource allocation and predictive visual analytics for community policing and law enforcement. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1863–1872, 2014.

[229] Miranda, F.; Doraiswamy, H.; Lage, M.; Zhao, K.; Goncalves, B.; Wilson, L.; Hsieh, M.; Silva, C. T. Urban pulse: Capturing the rhythm of cities. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 791–800, 2017.

[230] Purwantiningsih, O.; Sallaberry, A.; Andary, S.; Seilles, A.; Azfie, J. Visual analysis of body movement in serious games for healthcare. In: Proceedings of the IEEE Pacific Visualization Symposium, 229–233, 2016.

[231] Riehmann, P.; Kiesel, D.; Kohlhaas, M.; Froehlich, B. Visualizing a thinker's life. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 4, 1803–1816, 2019.

[232] Sacha, D.; Al-Masoudi, F.; Stein, M.; Schreck, T.; Keim, D. A.; Andrienko, G.; Janetzko, H. Dynamic visual abstraction of soccer movement. *Computer Graphics Forum* Vol. 36, No. 3, 305–315, 2017.

[233] Sarikaya, A.; Correli, M.; Dinis, J. M.; O'Connor, D. H.; Gleicher, M. Visualizing co-occurrence of events in populations of viral genome sequences. *Computer Graphics Forum* Vol. 35, No. 3, 151–160, 2016.

[234] Shi, C. L.; Wu, Y. C.; Liu, S. X.; Zhou, H.; Qu, H. M. LoyalTracker: Visualizing loyalty dynamics in search engines. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1733–1742, 2014.

[235] Steiger, M.; Bernard, J.; Mittelstädt, S.; Lücke-Tieke, H.; Keim, D.; May, T.; Kohlhammer, J. Visual analysis of time-series similarities for anomaly detection in sensor networks. *Computer Graphics Forum* Vol. 33, No. 3, 401–410, 2014.

[236] Stopar, L.; Skraba, P.; Grobelnik, M.; Mladenic, D. StreamStory: Exploring multivariate time series on multiple scales. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 4, 1788–1802, 2019.

[237] Sultanum, N.; Singh, D.; Brudno, M.; Chevalier, F. Doccurate: A curation-based approach for clinical text visualization. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 142–151, 2019.

[238] Sun, G. D.; Wu, Y. C.; Liu, S. X.; Peng, T. Q.; Zhu, J. J. H.; Liang, R. H. EvoRiver: Visual analysis of topic coopetition on social media. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1753–1762, 2014.

[239] Sung, C. Y.; Huang, X. Y.; Shen, Y. C.; Cherng, F. Y.; Lin, W. C.; Wang, H. C. Exploring online learners' interactive dynamics by visually analyzing their time-anchored comments. *Computer Graphics Forum* Vol. 36, No. 7, 145–155, 2017.

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

[241] Thom, D.; Kruger, R.; Ertl, T. Can twitter save lives? A broad-scale study on visual social media analytics for public safety. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 7, 1816–1829, 2016.

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

[243] Vehlow, C.; Beck, F.; Auwärter, P.; Weiskopf, D. Visualizing the evolution of communities in dynamic graphs. *Computer Graphics Forum* Vol. 34, No. 1, 277–288, 2015.

[244] Von Landesberger, T.; Brodkorb, F.; Roskosch, P.; Andrienko, N.; Andrienko, G.; Kerren, A. MobilityGraphs: Visual analysis of mass mobility dynamics via spatio-temporal graphs and clustering. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 11–20, 2016.

[245] Wang, X.; Dou, W.; Ma, Z.; Villalobos, J.; Chen, Y.; Kraft, T.; Ribarsky, W. I-SI: Scalable architecture for analyzing latent topical-level information from social media data. *Computer Graphics Forum* Vol. 31, No. 3, 1275–1284, 2012.

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

[247] Wang, Y.; Haleem, H.; Shi, C. L.; Wu, Y. H.; Zhao, X.; Fu, S. W.; Qu, H. Towards easy comparison of local businesses using online reviews. *Computer Graphics Forum* Vol. 37, No. 3, 63–74, 2018.

[248] Wei, F. R.; Liu, S. X.; Song, Y. Q.; Pan, S. M.; Zhou, M. X.; Qian, W. H.; Shi, L.; Tan, L.; Zhang, Q. TIARA: A visual exploratory text analytic system. In: Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 153–162, 2010.

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

[250] Wu, A. Y.; Qu, H. M. Multimodal analysis of video collections: Visual exploration of presentation techniques in TED talks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 7, 2429–2442, 2020.

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

[252] Wu, Y. C.; Chen, Z. T.; Sun, G. D.; Xie, X.; Cao, N.; Liu, S. X.; Cui, W. StreamExplorer: A multi-stage system for visually exploring events in social streams. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 10, 2758–2772, 2018.

[253] Wu, Y. C.; Liu, S. X.; Yan, K.; Liu, M. C.; Wu, F. Z. OpinionFlow: Visual analysis of opinion diffusion on social media. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1763–1772, 2014.

[254] Wu, Y. H.; Pitipornvivat, N.; Zhao, J.; Yang, S. X.; Huang, G. W.; Qu, H. M. egoSlider: Visual analysis of egocentric network evolution. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 1, 260–269, 2016.

[255] Xie, C.; Chen, W.; Huang, X. X.; Hu, Y. Q.; Barlowe, S.; Yang, J. VAET: A visual analytics approach for E-transactions time-series. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1743–1752, 2014.

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

[257] Xu, J.; Tao, Y. B.; Yan, Y. Y.; Lin, H. Exploring evolution of dynamic networks via diachronic node embeddings. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 7, 2387–2402, 2020.

[258] Xu, P. P.; Mei, H. H.; Ren, L.; Chen, W. ViDX: Visual diagnostics of assembly line performance in smart factories. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 1, 291–300, 2017.

[259] Xu, P. P.; Wu, Y. C.; Wei, E. X.; Peng, T. Q.; Liu, S. X.; Zhu, J. J.; Qu. H. Visual analysis of topic competition on social media. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 12, 2012–2021, 2013.

[260] Yu, L.; Wu, W.; Li, X.; Li, G.; Ng, W. S.; Ng, S.-K.; Huang, Z.; Arunan, A.; Watt, H. M. 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, 49–56, 2015.

[261] Garcia Zanabria, G.; Alvarenga Silveira, J.; Poco, J.; Paiva, A.; Batista Nery, M.; Silva, C. T.; de Abreu, S. F. A.; Nonato, L. G. CrimAnalyzer: Understanding crime patterns in São Paulo. *IEEE Transactions on Visualization and Computer Graphics* doi: 10.1109/TVCG.2019.2947515, 2019.

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

[263] Zeng, H. P.; Wang, X. B.; Wu, A. Y.; Wang, Y.; Li, Q.; Endert, A.; Qu, H. EmoCo: Visual analysis of emotion coherence in presentation videos. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 927–937, 2019.

[264] Zeng, W.; Fu, C. W.; Müller Arisona, S.; Erath, A.; Qu, H. Visualizing waypoints-constrained origin-destination patterns for massive transportation data. *Computer Graphics Forum* Vol. 35, No. 8, 95–107, 2016.

[265] Zhang, J. W.; Ahlbrand, B.; Malik, A.; Chae, J.; Min, Z. Y.; Ko, S.; Ebert, D. S. A visual analytics framework for microblog data analysis at multiple scales of aggregation. *Computer Graphics Forum* Vol. 35, No. 3, 441–450, 2016.

[266] Zhang, J. W.; E, Y. L.; Ma, J.; Zhao, Y. H.; Xu, B. H.; Sun, L. T.; Chen, J.; Yuan, X. Visual analysis of public utility service problems in a metropolis. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1843–1852, 2014.

[267] Zhao, J.; Cao, N.; Wen, Z.; Song, Y. L.; Lin, Y. R.; Collins, C. #FluxFlow: Visual analysis of anomalous information spreading on social media. *IEEE Transactions on Visualization and Computer Graphics* Vol. 20, No. 12, 1773–1782, 2014.

[268] Zhao, Y.; Luo, X. B.; Lin, X. R.; Wang, H. R.; Kui, X. Y.; Zhou, F. F.; Wang, J.; Chen, Y.; Chen, W. Visual analytics for electromagnetic situation awareness in radio monitoring and management. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 1, 590–600, 2020.

[269] Zhou, Z. G.; Meng, L. H.; Tang, C.; Zhao, Y.; Guo, Z. Y.; Hu, M. X.; Chen, W. Visual abstraction of large scale geospatial origin-destination movement data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 25, No. 1, 43–53, 2019.

[270] Zhou, Z. G.; Ye, Z. F.; Liu, Y. N.; Liu, F.; Tao, Y. B.; Su, W. H. Visual analytics for spatial clusters of air-quality data. *IEEE Computer Graphics and Applications* Vol. 37, No. 5, 98–105, 2017.

[271] Tian, T.; Zhu, J. Max-margin majority voting for learning from crowds. In: Proceedings of the Advances in Neural Information Processing Systems, 1621–1629, 2015.

[272] Ng, A. Machine learning and AI via brain simulations. 2013. Available at https://ai.stanford.edu/~ang/slides/DeepLearning-Mar2013.pptx.

[273] Nilsson, N. J. *Introduction to Machine Learning: An Early Draft of a Proposed Textbook*. 2005. Available at https://ai.stanford.edu/~nilsson/MLBOOK.pdf.

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

[275] Lee, K.; Lee, H.; Lee, K.; Shin, J. Training confidence-calibrated classifiers for detecting ut-of-distribution samples. *arXiv preprint* arXiv:1711.09325, 2018.

[276] Liu, M. C.; Jiang, L.; Liu, J. L.; Wang, X. T.; Zhu, J.; Liu, S. X. Improving learning-from-crowds through expert validation. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence, 2329–2336, 2017.

[277] Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; Fei-Fei, L. ImageNet

large scale visual recognition challenge. *International Journal of Computer Vision* Vol. 115, No. 3, 211–252, 2015.

[278] Chandrashekar, G.; Sahin, F. A survey on feature selection methods. *Computers & Electrical Engineering* Vol. 40, No. 1, 16–28, 2014.

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

[280] Tzeng, F.-Y.; Ma, K.-L. Opening the black box—Data driven visualization of neural networks. In: Proceedings of the IEEE Conference on Visualization, 383–390, 2005.

[281] Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M. et al. TensorFlow: Large-scale machine learning on heterogeneous distributed systems, *arXiv preprint* arXiv:1603.04467, 2015.

[282] Ming, Y.; Xu, P. P.; Qu, H. M.; Ren, L. Interpretable and steerable sequence learning via prototypes. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 903–913, 2019.

[283] Liu, S. X.; Cui, W. W.; Wu, Y. C.; Liu, M. C. A survey on information visualization: Recent advances and challenges. *The Visual Computer* Vol. 30, No. 12, 1373–1393, 2014.

[284] Ma, Z.; Dou, W.; Wang, X.; Akella, S. 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, 260–267, 2013.

[285] Kosara, R.; Bendix, F.; Hauser, H. Parallel sets: Interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics* Vol. 12, No. 4, 558–568, 2006.

[286] Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; Dean, J. Distributed representations of words and phrases and their compositionality. In: Proceedings of the Advances in Neural Information Processing Systems, 3111–3119, 2013.

[287] Blei, D. M.; Ng, A. Y.; Jordan, M. I. Latent Dirichlet allocation. *Journal of Machine Learning Research* Vol. 3, 993–1022, 2003.

[288] Teh, Y. W.; Jordan, M. I.; Beal, M. J.; Blei, D. M. Hierarchical dirichlet processes. *Journal of the American Statistical Association* Vol. 101, No. 476, 1566–1581, 2006.

[289] Wang, X. T.; Liu, S. X.; Song, Y. Q.; Guo, B. N. Mining evolutionary multi-branch trees from text streams. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 722–730, 2013.

[290] Li, Y. F.; Guo, L. Z.; Zhou, Z. H. Towards safe weakly supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* doi: 10.1109/TPAMI.2019.2922396, 2019.

[291] Li, Y. F.; Wang, S. B.; Zhou, Z. H. Graph quality judgement: A large margin expedition. In: Proceedings of the International Joint Conference on Artificial Intelligence, 1725–1731, 2016.

[292] Zhou, Z. H. A brief introduction to weakly supervised learning. *National Science Review* Vol. 5, No. 1, 44–53, 2018.

[293] Foulds, J.; Frank, E. A review of multi-instance learning assumptions. *The Knowledge Engineering Review* Vol. 25, No. 1, 1–25, 2010.

[294] Zhou, Z. H. Multi-instance learning from supervised view. *Journal of Computer Science and Technology* Vol. 21, No. 5, 800–809, 2006.

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

[296] Wang, Q. W.; Yuan, J.; Chen, S. X.; Su, H.; Qu, H. M.; Liu, S. X. Visual genealogy of deep neural networks. *IEEE Transactions on Visualization and Computer Graphics* Vol. 26, No. 11, 3340–3352, 2020.

[297] Ayinde, B. O.; Zurada, J. M. Building eficient ConvNets using redundant feature pruning. *arXiv preprint* arXiv:1802.07653, 2018.

[298] Baltrusaitis, T.; Ahuja, C.; Morency, L. P. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 41, No. 2, 423–443, 2019.

[299] Lu, J.; Batra, D.; Parikh, D.; Lee, S. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: Proceedings of the Advances in Neural Information Processing Systems, 13–23, 2019.

[300] Lu, J.; Liu, A. J.; Dong, F.; Gu, F.; Gama, J.; Zhang, G. Q. Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering* Vol. 31, No. 12, 2346–2363, 2018.

[301] Yang, W.; Li, Z.; Liu, M.; Lu, Y.; Cao, K.; Maciejewski, R.; Liu, S. Diagnosing concept drift with visual analytics. *arXiv preprint* arXiv:2007.14372, 2020.

[302] Wang, X.; Chen, W.; Xia, J.; Chen, Z.; Xu, D.; Wu, X.; Xu, M.; Schreck, T. Conceptexplorer: Visual analysis of concept drifts in multi-source time-series data. *arXiv preprint* arXiv:2007.15272, 2020.

[303] Liu, S.; Andrienko, G.; Wu, Y.; Cao, N.; Jiang, L.; Shi, C.; Wang, Y.-S.; Hong, S. Steering data quality with visual analytics: The complexity challenge. *Visual Informatics* Vol. 2, No. 4, 191–197, 2018.

Jun Yuan is currently a Ph.D. student at Tsinghua University. His research interests are in explainable artificial intelligence. He received his 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 his B.S. degree from the University of Science and Technology of China.

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

Mengchen Liu is a senior researcher at Microsoft. His research interests include explainable AI and computer vision. He received his B.S. degree in electronics engineering and his Ph.D. degree in computer science from Tsinghua University. He has served as a PC member and reviewer for various conferences and journals.

Jiazhi Xia is an associate professor in the School of Computer Science and Engineering at Central South University. He received his Ph.D. degree in computer science from Nanyang Technological University, Singapore in 2011 and obtained his M.S. and B.S. degrees in computer science and technology from Zhejiang University in 2008 and 2005, respectively. His research interests include 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 has worked as a research staff member at IBM China Research Lab and a lead researcher at Microsoft Research Asia. She received her B.S. and M.S. degree from Harbin Institute of Technology, and her Ph.D. degree from Tsinghua University. She is an Associate Editor-in-Chief of *IEEE Trans. Vis. Comput. Graph.*