DOI: 10.1111/cgf.14859
EUROVIS 2023
S. Bruckner, R. Raidou, and C. Turkay
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 42 (2023), Number 3
STAR – State of The Art Report

VA + Embeddings STAR: A State-of-the-Art Report on the Use of Embeddings in Visual Analytics

Z. Huang[1], D. Witschard[2], K. Kucher[1], and A. Kerren[1,2]

[1]Department of Science and Technology, Linköping University, Sweden
[2]Department of Computer Science and Media Technology, Linnaeus University, Sweden

Abstract

Over the past years, an increasing number of publications in information visualization, especially within the field of visual analytics, have mentioned the term "embedding" when describing the computational approach. Within this context, embeddings are usually (relatively) low-dimensional, distributed representations of various data types (such as texts or graphs), and since they have proven to be extremely useful for a variety of data analysis tasks across various disciplines and fields, they have become widely used. Existing visualization approaches aim to either support exploration and interpretation of the embedding space through visual representation and interaction, or aim to use embeddings as part of the computational pipeline for addressing downstream analytical tasks. To the best of our knowledge, this is the first survey that takes a detailed look at embedding methods through the lens of visual analytics, and the purpose of our survey article is to provide a systematic overview of the state of the art within the emerging field of embedding visualization. We design a categorization scheme for our approach, analyze the current research frontier based on peer-reviewed publications, and discuss existing trends, challenges, and potential research directions for using embeddings in the context of visual analytics. Furthermore, we provide an interactive survey browser for the collected and categorized survey data, which currently includes 122 entries that appeared between 2007 and 2023.

Keywords: embedding techniques, distributed representations, visual analytics, visualization

ACM CCS: • Human-centered computing → Visual analytics; • Human-centered computing → Information visualization; • Human-centered computing → Visualization systems and tools; • Computing methodologies → Machine learning; • Applied computing

1. Introduction

In recent years, the concept of *embedding technology* has gained a lot of attention within the research community. The term "embedding" refers to the process and/or the results of projecting data entities into a (relatively) low-dimensional space where the similarities of different points reflect their semantic closeness in the original space [MLMRC18, GWW19]. (Besides the process, the resulting *embedding vectors* are often referred to as "embeddings", too.) Various embedding approaches have been developed based on a target application domain and data types, such as numerical values [BCV13], word/text [MSC*13, LM14, AX19, FYC*22], graphs/networks [HYL17, GF18, ZYZZ20], media data [AGH*23], or a combination of those [BCV13, WMWG17]. Learning effective embeddings is crucial for performing downstream tasks in various fields, such as information retrieval [JSR*19], natural language processing [EAKC*20], social network analysis [AAM*21], and urban planning [MHL*20]. To further enhance the accuracy, explainability, and credibility of either the embedding process or the higher-level objectives, many studies incorporate a human-in-the-loop process in addition to optimizing the computational components [EHR*14]. There is a need to explore relationships within the embedding space and interpret the embedding process—and that is where visual analytics steps in.

In information visualization (InfoVis) and visual analytics (VA), many studies have been carried out on creating interactive user interfaces (UI) to understand embeddings over the past years. The rapid development of embedding methods, especially in deep representation learning [GWW19], enables many VA tools that acquire non-numerical data to perform new types of joint analysis, creating customized visual encodings and interactions. As Figure 1 shows, embeddings can be computed as part of VA pipelines or provided as part of input data. They are associated with various approaches for performing analytical tasks, encoding visual meanings, and enabling user interactions within the embedding space. These embeddings can be evaluated with regard to both their computational processes and their visual analytical aspects. The methodology and application scenarios for embedding technologies are highly heterogeneous. Furthermore, the terminology of the field is far from being standardized [KØSV18], making it hard to get a comprehensive overview for the respective interested readers.

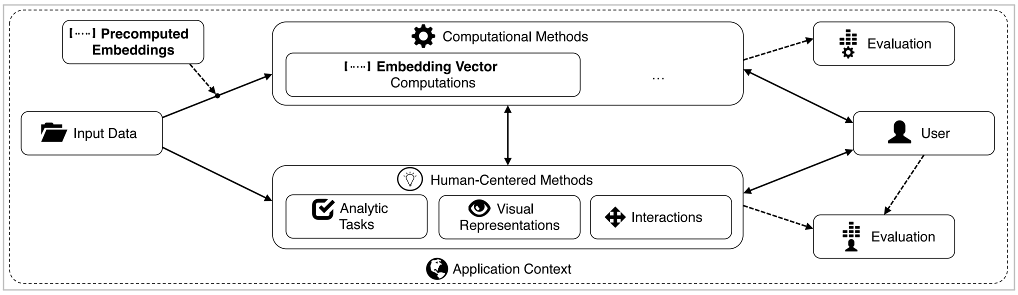

Figure 1: *A generic visual analytic pipeline instantiated for the VA approaches and tools that either include embedding computations as part of the computational approach, or make use of externally computed data embeddings. The aspects such as the overall domain application context, target user, input data, resulting embedding vectors, computational and human-centered methods, and corresponding evaluation concerns define the scope of our survey and the proposed categorization discussed throughout this manuscript.*

For the research community in visualization, understanding how embeddings play a role here is important yet challenging, and the reason is twofold. First, in line with the rapid advances in embedding technologies, there is an increasing number of InfoVis and VA papers mentioning embeddings as part of their computational pipeline. Second, due to the heterogeneous nature of this topic, there are various domain-specific visualization approaches and representations. Similarities and differences between the domains might provide insight into the common ground and gaps for future research. They also act as success stories of applying VA that might attract further interest from researchers and practitioners in other disciplines and domains. However, there is currently a lack of a systematic taxonomy and overview of the fields of InfoVis and VA that would offer a comprehensive examination of how embedding approaches are used within these fields.

This survey attempts to address this gap, specifically within the context of information visualization and visual analytics. Furthermore, we restrict its scope to only focus on peer-reviewed publications which meet at least one of the following two criteria: (1) embeddings are explicitly represented in at least one visual representation, and (2) embeddings of some source data items are explicitly computed within the described pipeline. We also limit our scope to embedding spaces larger than 4 dimensions (4D) only. This is done to exclude a potentially much broader selection of VA papers that use general dimensionality reduction (DR) techniques [LMW*17, SZS*17, EMK*21] solely to compute a layout in 2D/3D/3D+time.

Figure 1 illustrates the basic aspects and challenges of the use of embeddings mentioned above in the context of a typical visual analytic pipeline inspired by the works by Keim et al. [KAF*08] and Sacha et al. [SSS*14]. A visual analytic approach is typically designed to address a particular problem, be it a particular domain application or a more generic, domain-agnostic problem. In order to address the needs of the respective target users, the input data is fed into computational and visual/interactive modules, which are also communicating with each other, e.g., to provide visual representations of computational results. Data embeddings might be computed by such VA approaches as part of the computational module and used as an intermediate representation for further analyses (e.g., computing word vectors for further document clustering), or alternatively, fed to the visualization module (e.g., to provide a visual summary of the embedding vector values). Furthermore, the embeddings computed outside of the VA tool itself might be provided as part of the input data. The precise focus of this survey thus lies on the VA approaches that explicitly make use of data embeddings, according to the two criteria named above (explicit visual representation and/or explicit computation); and for this reason, we also consciously limit the discussion of potentially numerous coordinated views and interactions [Rob07] available in complex VA applications beyond the ones directly relevant to data embeddings.

The contributions of our work are as follows:

- We present a categorization schema for VA approaches involving embeddings, which takes not only visualization-related aspects but also the computational aspects and domain applications into account.
- Based on the proposed categorization and our corpus containing 122 papers, we identified 15 subfields (scattered within 9 domains) that use embedding techniques in their VA systems. For each subfield, we provide a collection of state-of-the-art papers and an overview of their use of embeddings integrated into visualization systems, including motivations, analytical tasks, and VA-related commonalities.
- Besides within-field overviews, we offer a cross-domain summary and discussion of existing trends, challenges, and potential research directions from two perspectives: using embedding techniques to enhance a VA system (embeddings4VA) and using visual interfaces to understand and explore embedding spaces (VA4embeddings).
- We provide an interactive survey browser for the collected and categorized survey data, available online at

  https://va-embeddings-browser.ivis.itn.liu.se/

1.1. How to Use This Survey

Our article aims to incentivize the visualization community to reflect and focus on this important topic directly (rather than approaching it purely from the ML or DR perspectives) and, to the best of our knowledge, this is the first survey that takes a detailed look at embedding methods through the lens of VA. We propose the following usage scenarios:

- read as a general overview, for any researcher who needs a comprehensive summary of how embedding data is currently being visualized in different VA tools but who does not have this as their main domain of research;
- use the survey browser to identify visualization techniques that meet specific criteria, for researchers who are searching for inspiring examples/solutions to specific problems within their domain of expertise;
- read to identify gaps and/or research opportunities, for any researcher who intends to focus more specifically on embedding visualization within their field of research; and
- use the supplementary material to dive into details on a particular category, frequent category pattern, or co-authorship network.

The rest of this manuscript is organized as follows: in Section 2, we discuss the terminology, origins, and variety of the existing embedding computation approaches in relation to the topic and scope of our survey. Section 3 positions our work with respect to the previous survey articles within and beyond the visualization field. Next, we discuss the methodology of data collection and annotation in Section 4. Section 5 focuses on the resulting categorization of VA + embedding approaches: here, we discuss various aspects, from the application domain to the details of computational and visual/interactive components evaluation, and provide examples of the respective approaches. Besides focusing on the results and examples for each particular category, we carry out further data-driven analyses in Section 6, including temporal, category correlation, topical, and (co-)authorship analyses. We reflect on and discuss the findings and limitations of this study in Section 7, as well as the open challenges identified in this field. Finally, Section 8 concludes this manuscript.

2. Background

Before going into the details of the survey results, we first want to provide some general background information regarding embeddings and embedding technologies. First, we note that the word *embedding* can have different meanings in the English language and that there, to the best of our knowledge, is no formally accepted definition that is applicable to our selected scope. Therefore, we start by noting that, within the scope of the surveyed corpus, the word *embedding* is generally used in the following meaning: *an alternative representation of some underlying data which aims to preserve important characteristics and/or relations of the data points* [CTL18,HPX*21,WJM*22,YHZ22]. Consequently, *embedding technology* is used in the meaning: *a method that transforms data points into a different form while preserving important characteristics and/or relations*. Precisely what these important characteristics/relations are may vary from case to case. Still, in general, the main goal of embedding algorithms is to preserve the closeness/similarity of the data points in the original data space, so that points that are close/similar end up being close/similar also in the embedding space (which is often referred to as the *latent space*). We also note that for all but one of the surveyed publications, the embedding representation is in the form of a numerical vector—so the words *embedding* and *vector* are often used interchangeably (although the single exception should serve as a reminder that this does not always need to be the case). For the reader seeking a more compact wording, we offer the following: "in most cases, embedding algorithms represent domain data as numerical vectors that preserve similarity (from the original data space) and allow for efficient computational analysis."

The origin of embedding technologies arguably lies in the field of natural language processing (NLP) and computational linguistics (CL), where several different technologies (such as statistical, matrix-based, and machine-learning models) have been developed to transform the words of a corpus into numeric vector representations which preserve the semantic relationships between them [BDVJ03, CWB*11]. The primary rationale for developing such types of transformations is that it is much easier to calculate the distance between two vectors than computationally deriving the semantic similarity of two words directly from their textual representations (i.e., the level of semantic closeness/similarity between two words is instead established by a calculation using the numerical embedding vectors as proxies). A more detailed explanation of the development, and inner workings, of different word/text embedding technologies, falls outside the scope of this survey [OMK21]. Therefore, we would like to note that, with the arrival of deep learning (DL) models and algorithms such as word2vec [MSC*13] and bidirectional encoder representations from transformers (BERT) [LM14,DCLT19], the initial hopes have been more than exceeded in the sense that the complex semantic relationships between words and/or text sequences can be captured by such embedding algorithms and thereafter exploited by vector calculations.

After the initial success of embedding technologies within the field of NLP, attempts were made to adapt the general idea to other types of data suffering from the same type of inherent challenges as texts (i.e., the seemingly unstructured and complex nature of the original data making it hard to perform direct calculations on it). A prominent example of such a data type is graph/network data since it suffers from the fact that many topology-related calculations are cumbersome and computationally expensive to perform. However, with the discovery that the core ideas of the word2vec learning model could be applied to network topology data (as in the node2vec algorithm [GL16], which, loosely speaking, treats nodes as "words" and random walks in the network as "sentences"), this field also became an important area of development for embedding algorithms. From the initial methods for pure graph embedding [GF18] (supporting tasks such as graph topology comparison and subgraph search), the field has evolved towards algorithms that also take into account the attributed data on nodes and edges [CWPZ19], which in turn allows for even more complex analysis scenarios, e.g., similarity comparisons of nodes jointly evaluated on topological position and attribute resemblance.

The concept of *dimensionality* is important within the field of embeddings since it puts a restriction on how much of the desired characteristics of the underlying data the algorithm will be able to capture. All other things being equal, a vector of higher dimension (i.e., with greater length) will have higher "expressiveness" than a vector of lower dimensions (i.e., with shorter length), simply because there are more possible ways to group data points (i.e., to express their relations) in a larger embedding space than in a smaller one. On the other hand, a longer vector will introduce an overhead, both in computational load and in memory usage, so a trade-off is usually needed. Furthermore, it is not uncommon to see statements indicating that embeddings are "low-dimensional" vector representations of the underlying data, which may seem contradictory to the fact that many such algorithms use several hundred dimensions for their yielded vectors. Whether such a representation of a word in a text, or a node of a network, really can be seen as truly low-dimensional is debatable. However, with the knowledge that complete texts or networks can be captured by such vectors, the notion of (relatively) low-dimensional becomes more apparent.

We would like to end this section by noting that embedding technologies are a *relatively* new addition to the ML/AI field [BCV13], and as such, a lot of research effort has been directed toward developing and using rather than visualizing. However, in accordance with the large trend of eXplainable AI (XAI), several initiatives have proven that there is much to be learned from opening the "black box" of embeddings and embedding spaces. With this survey, we hope to provide the reader with a structured overview of this emerging field. We also hope that it can catalyze further advances in developing visualizations for embedding data.

3. Related Surveys

Existing surveys published in other disciplines and fields focus mainly on exploring the computational aspect of embedding techniques, while visualization is only mentioned in passing, typically in relation to static plots with limited interaction [WMWG17, GF18, CWPZ19, HRZ*20]. For instance, Cui et al. [CWPZ19] present a network embedding survey focusing on challenges and opportunities in this field, as well as relationships across different methods. The article includes one paragraph on network visualization, but the authors frame it as a visual demonstration of how embeddings could preserve the intrinsic structure of the network. As for most surveys on embeddings in the field of NLP, visualization is barely mentioned [KØSV18, LY18, LKB20].

We have also examined recent visualization surveys that mention embedding techniques. The topics of those surveys range from visualization techniques for explainable machine learning [HKPC18, CMJK20, CMJ*20, SEAG*21] to specific types of networks and maps [HHS20, WNT*20] or text data [KPK18, LWC*19]. There are also several surveys analyzing how specific computational techniques benefit VA by focusing on integrating general ML methods [ERT*17], as well as DR algorithms [LMW*17, SXG*19, EMK*21]. Nonato et al. [NA19] focus on DR projections for multidimensional data, specifically on handling distortions (such as false or missing neighbors), so that the negative impact on the visual analysis of the projection can be kept limited.

Furthermore, we have identified two surveys within InfoVis and VA that come closest to our survey scope. First, Sacha et al. [SZS*17] focus on how people can interact with the DR pipeline and how visualization can be integrated into an interactive dimensionality reduction process. The methodology used in this survey is similar to ours. However, they concentrate mainly on DR, which typically maps data from high dimensions to 2D or 3D. This mapping often has implications for the design of the complete workflow and any (interactive) visualizations. While DR and embedding methods are deeply intertwined, our survey holds a different emphasis and perspective: we focus on the techniques and application scenarios involving embedded data representations in spaces higher than 4 dimensions, thus separating our focus from typical DR use in visualization. Second, Wu et al. [WWS*21] have recently provided a survey on how AI applications could benefit data visualization. Some of the techniques for exploring the embedding space within an ML model are discussed in their paper; however, they are mentioned only as part of the broad coverage of their survey, with limited analysis and no comparisons of those particular methods across domains.

4. Methodology

The methodology for conducting this survey is based on a number of previous visualization surveys and meta-analyses [CK15, KK15, KKLS17, KPK18, KMK18, KK19, CMJK20, CMJ*20, WJMK21], and further inspired by other surveys and sensemaking models [PC05, BKW16, SEAG*21].

Figure 2 provides an overview of the activities and outcomes of our work on this survey. The initial stages of this process were dedicated to the definition of scope and inclusion/exclusion criteria for publications that act as the basis for survey entries (we should note that InfoVis and VA tools, approaches, and techniques are the primary units of our survey rather than publications themselves).

The initial pilot collection of candidate publications was compiled by querying and manually screening articles published in two journals, Computer Graphics Forum (CGF) and IEEE Transactions on Visualization and Computer Graphics (TVCG). We arrived at 425 hits based on several search queries that contain the keyword "embedding" and further keywords such as "visualization", "interface", "interact", and derivatives with similar meanings. Initial manual screening of the publication titles and abstracts was conducted to ensure that the respective publications were related to interactive visualization and embedding approaches (at this point, with a very broad notion of "embedding"). After this stage, the resulting pilot data set included 69 articles.

As we did not consider the basic keyword search to be sufficient, the next step involved a closer inspection of the pilot data set by 3 annotators with the goal of marking each paper as relevant or irrelevant to the scope of our survey. At this stage, we formulated the inclusion/exclusion criteria and applied them to the respective candidate papers (based on title and abstract, as the procedure was to be applied to a large set of candidates afterwards). The resulting annotations from individual annotators were afterwards compared and discussed in order to resolve conflicts and clarify the criteria. At this step, Cohen's kappa values were calculated to estimate the pairwise inter-annotator reliability [AP08], with the values of 0.605, 0.593, and 0.673 for three pairs of annotators from among the authors of this survey. These values indicate moderate to substantial, albeit far from perfect agreement [AP08]; thus, the particular disagreement cases were discussed by the annotators, and furthermore, the decision to check the full contents of candidate papers in case of doubt was made (which is not practical for clear inclusion/exclusion cases).

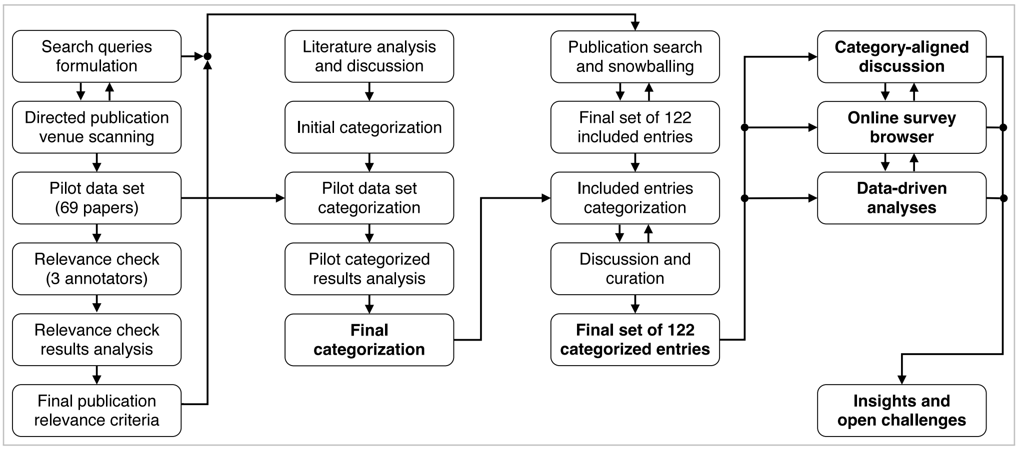

Figure 2: *Overview of the methodology and data involved in the preparation of this state-of-the-art report. The main outcomes and contributions of our work are highlighted in bold.*

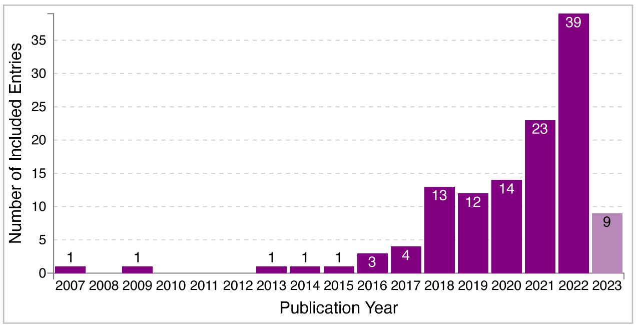

Figure 3: *Temporal distribution of the entries included in our survey (**122** in total as of April 2023) based on the base publication year.*

While the overall inclusion/exclusion criteria stated as part of the survey scope in Section 1 proved to be adequate, such detailed checks were helpful in order to exclude, for instance, papers that did not provide evidence of embedding use beyond 2D/3D DR projections. Some of the interesting borderline cases that we eventually removed from the survey set include, for instance, the work by Raidou et al. [RvdHD*15] that explicitly mentions embedding high-dimensional data points into an abstract 2D space. Similar considerations led us to exclude the papers by Sohns et al. [SGL22], Eckelt et al. [EHA*22], and Zeng et al. [ZZL*22], which all present valuable contributions, but do not fit the scope of this survey.

The next major step consisted of the categorization design for the included survey entries. While the resulting categorization is presented in detail in the following section, it is worth noting that in order to refine the design, an initial version of categorization was used for annotating the included papers from the pilot set described above. We iteratively discussed the results and ambiguities in order to finalize the categorization schema. Then, the search for further candidate papers (involving the search in particular journals/libraries, but also further Google Scholar search and *snowballing* [Woh14] to extend the candidate set), relevance check with inclusion/exclusion criteria, and annotation according to the categorization took place.

The resulting set of **122** entries is summarized with respect to the underlying publication year in Figure 3. These entries are based on articles and papers from the broad field of visualization (IEEE TVCG, CGF, IEEE VIS, EuroVis, PacificVis, Inf Vis, JoV, and other venues), but also a number of publications from other fields such as NLP, for instance (with publications from several ACL venues). Rather than limiting our survey to a particular group of venues, we have been interested in discovering the various domain applications of visual analytics related to the use of embeddings, and in the future we intend to keep extending the survey data set using the approach discussed below; further candidate entries are available for annotation, and, as Figure 3 demonstrates, the past several years have resulted in a number of relevant publications indicating interest in this topic, so we expect to see a lot of interesting contributions in the near future.

To facilitate data exploration and present the results, we have developed (and extensively used ourselves) an interactive survey browser demonstrated in Figure 4. The overall user interface design here follows the existing survey browsers [KK15,KPK18,CMJ*20] and allows for exploration of individual survey entries (the grid of thumbnails on the right) as well as faceted search (the filters on the left). The implementation was, however, extended to accommodate the free-text detail categories and potential multiple references and URLs specified per survey entry. For instance, while the entry for DRIFT by Pocco et al. is based on the 2022 journal article [PdSP*22] identified during the literature search stage, another reference for the related conference paper [PPV*21] is also mentioned in the browser entry. This might come across as a trivial implementation detail, but the aim here is to keep supporting and extending the survey data in the future for the benefit of the visualization research community, while promoting the cases of several successful applications of visual analytic approaches with embeddings that the respective authors might describe in a series of publications. Additional URLs included for the survey entries can

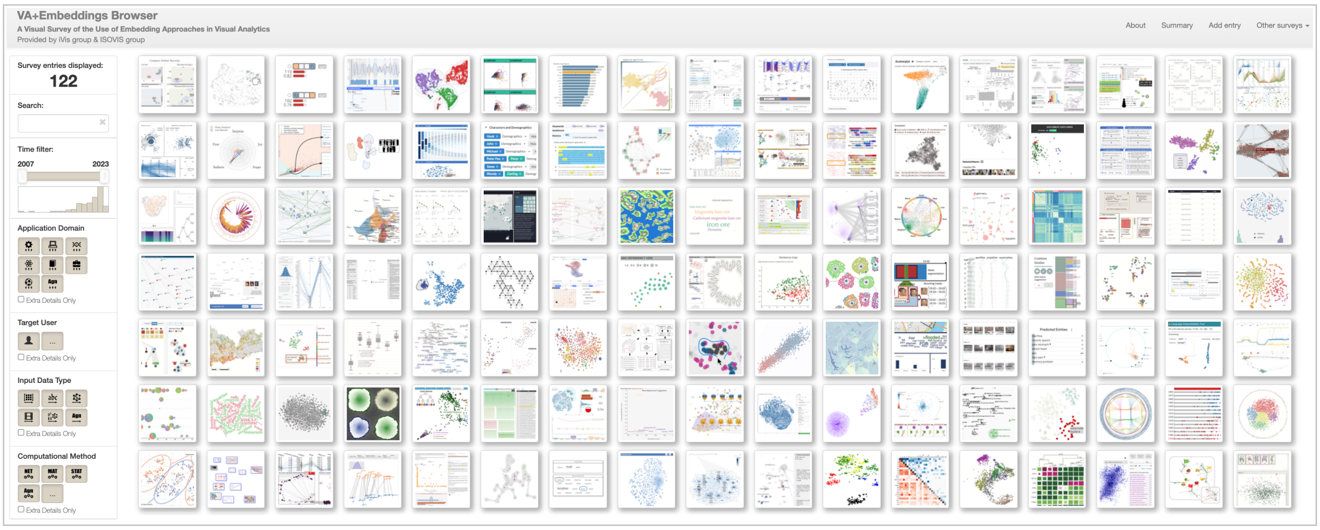

Figure 4: *The online survey browser accompanying this manuscript.*

also typically lead to source code repositories, project websites, or online demos, thus contributing to the visibility of the respective research.

As indicated in Figure 2, the remaining steps of the work on this survey included the annotation and analysis of publications with respect to corresponding categories (performed mainly by the first two authors of this survey, and presented in the next section), further data-driven analyses of the collected survey data (see Section 6), and discussion of the respective outcomes (Section 7).

5. Categorization of VA + Embedding Approaches

As presented in Figure 1 and discussed in Section 1, the categorization design for this survey aims to ensure that a wide range of aspects, including application domains, user profile, embedding computation methods, and visual/interactive design concerns, are addressed. Table 1 presents the resulting categorization, with the previously mentioned top-level aspects comprising particular nested categories. Similar to the previous surveys [KK15, KPK18, CMJ*20], the categories are generally not exclusive within the same group/aspect, e.g., a particular visual analytic approach could support several application domains. Furthermore, while the current survey data set includes **122** entries, certain aspects do not cover 100% of the entries. Additionally, as we had a pragmatic intention to keep the size of categorization constrained, not every possible computational model, visual representation, or evaluation approach could be promoted to a category of its own. Thus, we introduced additional free-text comment categories for most of the top-level aspects of this categorization. The notes made by individual annotators that were recorded within such free-text comments are available via the online survey browser. The counts of entries that include such notes are also provided in Table 1.

This section outlines the motivation for including particular categories in our proposed categorization, discusses their support within the current survey data set, highlights noteworthy findings, and describes the prominent examples.

5.1. Application Domain

Here, we present our categorization of application domains, explain the included subfields, and provide an overview of papers in our corpus for each subfield within each categorized domain.

VA tools that use embedding technology are designed for a wide range of topics, such as enhancing ML model understanding [PDD*22], diagnosing graph embeddings [RSL*22], analyzing medical records [CEBV22], urban soundscapes [RMH*22], and product reviews [JCSM22]. Papers in different fields may contain similar or different initiatives, themes, and perspectives related to VA+embeddings. Classifying papers according to their application domains is necessary to gain insights into within-field applications and make between-field comparisons. This section presents our categorization of application domains, which serves as the basis for our analysis and further discussion in Section 7. As shown in Table 1, approximately half of the publications in our corpus were classified within the domain **ML, AI, data science** (57). This distribution is consistent with the common mentioning of the word "embeddings" under a machine learning context for VA-assisted ML (VIS4ML) applications [SZS*17]. The remaining publications were distributed across several other domains, including **computing** (24), **humanities, social sciences, and education** (23), **life sciences and medicine** (13), **domain-agnostic** (12), **business, management, governance, law** (12), **physical sciences, engineering, and mathematics** (3), and **sports and entertainment** (3).

We choose the above domain categorization with two main considerations: practicality and representativeness. Concerning the number of publications surveyed and the wide variety in the context of each paper, a too fine-grained categorization could risk making the overview less comprehensive. Thus, we opted for fewer

Table 1: *The overview of the categorized survey data set. Each row contains the number of corresponding techniques/entries in our data set as of April 21, 2023. The percentage relative to the current total of **122** techniques is also illustrated by heatmap-style icons (note that the categories within each group/aspect are not necessarily mutually exclusive).*

Application Domain	**122**
ML, AI, Data Science	57
Computing	24
Life Sciences and Medicine	13
Physical Sciences, Engineering, Mathematics	4
Humanities, Social Sciences, Education	23
Business, Management, Governance, Law	12
Sports and Entertainment	3
Domain-Agnostic	12
Free-text details	*122*

Target User	**32**
Explicit Target User Description	32
Free-text details	*33*

Input Data Type	**122**
Numerical/Tabular	25
Textual	70
Graph/Network	22
Image/Video/Audio	17
Mixed	17
Data-Agnostic	6
Free-text details	*59*

Computational Method	**116**
Neural Network	88
Matrix Analysis	33
Statistical Analysis	16
Computational Method-Agnostic	12
Free-text details	*122*

Embedding Vector Dimensionality	**122**
Under 50 Dims.	15
Between 50–500 Dims.	31
Above 500 Dims.	14
Embedding Dims. N/A	70

Visual Analytic Task	**118**
Model Construction	14
Model Debugging/Quality/Bias Control	8
Model Results Representation	63
Model Results Explanation	23
Interactive Exploration	53
Comparison/Selection	52
Free-text details	*122*

Visualization Aspects	**102**
Explicit Embedding Representation	102
Free-text details	*107*

Interaction Aspects	**77**
Interaction Techniques Support	77
Free-text details	*79*

Evaluation Aspects	**98**
Evaluation of Computational Components	53
Evaluation of Visual/Interactive Components	84
Free-text details	*101*

and broader categories in our categorization approach. To establish our domain categorization, we drew upon multiple sources, including the ACM Computing Classification System (CCS) and the list of application domains used by the IEEE Visualization and Graphics Technical Community (VGTC) and IEEE Visualization (IEEE VIS). In addition, we gathered insights from our previous research related to the trustworthiness of machine learning models (TrustML) visualization and NLP-related sentiment visualization [KPK18, CMJK20], and also from other surveys on visualization and embedding techniques [SKKC19, CWPZ19, LKB20, XZL*22].

The interdisciplinary nature of some publications makes them fit under more than one category. For instance, a VA tool that uses advanced interactive recurrent neural networks to analyze medical records [KCK*19] would fall under both the categories of **ML, AI, data science**, as well as **life sciences and medicine**. In borderline cases, we consider additional factors, such as the framing of a paper to make a final judgment.

While we preferred to avoid adding another level of nesting directly to our overall categorization, it is worth noting that we were able to identify some smaller groups of interesting examples within specific application domains. For example, medical records analysis [GFL*20] and mass cytometry [HPvU*16] (see Figure 5(a)) are put under the umbrella of **life sciences and medicine**, while **business, management, governance, law** includes examples from the fields such as market analysis [JCSM22], collaboration tools [XBL*18], and maintenance [ZFC*21] (see Figure 5(b)).

Thus, readers interested in using embedding approaches for a specific domain application area of visualization could navigate directly to the following paragraphs:

- Visualization for NLP and CL (Sect. 5.1.1)
- Neural Network Interpretation (Sect. 5.1.1)
- Performance and Software Visualization (Sect. 5.1.2)
- Visual Search (Sect. 5.1.2)
- Biological Data Visualization (Sect. 5.1.3)
- Visualization for Healthcare (Sect. 5.1.3)
- Social Media Visual Analytics (Sect. 5.1.4)
- Urban Visual Analytics (Sect. 5.1.4)
- Visualization for Public Safety (Sect. 5.1.4)
- Visualization for Traffic Flow (Sect. 5.1.4)
- Geo-text Data Visualization (Sect. 5.1.4)
- Visual Interfaces for Collaboration (Sect. 5.1.5)
- Visual Browsing of Multimedia Collections (Sect. 5.1.6)
- Visualization for Sports and Entertainment (Sect. 5.1.7)
- Visualization for Embedding Interpretation (Sect. 5.1.8)

5.1.1. ML, AI, Data Science

Exploring latent representations [FZCM20], decision boundaries [SGR*20], and network characteristics [JLB22] are essential to examine and understand machine learning models. Particularly as certain aspects of the embedding process or embedding space are considered research questions, visual analytics is an effective way to help ML understanding, diagnosis, and refinement [LWLZ17]. 57 papers in our corpus fit under the domain of *ML, AI, and Data Science*. Among them, we included 39 papers within the field of natural language processing and computational linguistics.

Visualization for NLP and CL Creating and using embeddings in NLP and CL is crucial for representing and capturing the context and content of words, phrases, sentences, and documents. VA + embedding techniques in this set focus on four themes: exploring the semantics and contextualization of embedding spaces [CTL18, LBT*18, EAKC*20, MWZ19, SSKEA21, GHM21, BN21, BCS22, VMZL22, LWZ*23, MM23], active learning and interpretation for language models [LCSEK19, TWB*20, SH20, ARCL21, LXW*21, SKB*21, SCR*23], data-driven information retrieval [CWDH09, BMS17, ZSHL18, KOK*18, DMdO19, RSBV21, PdSP*22, JWC*23], and annotation tools [SJB*17, BNL*18, PKL*18, MWJ22].

Interpreting embeddings at a word level is still the predominant topic. Instead of calculating static metrics, VA makes interactive comparisons across groups possible. For instance, *Embedding Comparator* by Boggust et al. creates visualizations to compare and link different local neighborhoods of words with the global embedding structure (Figure 5(c)) [BCS22]. For different word and concept groups, one may observe patterns of intersectional biases along different social categorizations [GHM21], structures for semantic and syntactic analogies [LBT*18], regularization processes in different dimensions [LWZ*23], and diachronic changes [CTL18].

Pre-trained language models are widely used to consider a word's context and move beyond probing classifiers. They are trained on a vast amount of text data, with multiple layers capturing different levels of information. VA tools in this set also align with the computational trends in visually exploring contextualized word embeddings. Unlike traditional approaches that rely on discovering distribution patterns and frequencies of keyword clusters, encoding contextualization involves using neural representations from pre-trained language models to convey the context of a word. Besides exploring a single embedding space, VA tools for understanding contextualized word embedding focused more on intra-layer properties and inter-layer differences [SKB*21]. For instance, Sevastjanova et al. [SKB*22] compare transformer-based language models' layerwise context specificities to help explain how contextualization propagates through BERT-like models (Figure 6(a)).

A visual approach of encoding contextual information is to add Voronoi maps of keywords [VMZL22] and colored glyphs [SCR*23] to partition the semantic space. However, it is challenging to deal with uncertainties arising from the projection of the dataset and the partitioning process. Although the common assumption is to ensure the resulting partitioning is non-overlapping [SSJ*22], there will be probabilities associated with each resulting set/area. As a result, in *KeywordSpace* [VMZL22], uncertainties are represented by glyph coloring. Alternatively, Sohns et al. propose the use of specific scalar-field projections and geometric set-based visualizations for such spatial augmentations to handle uncertainties [SSJ*22].

Contextualizing word embeddings offers opportunities to provide guidance at different levels. Many existing VA tools for exploring word embeddings allow comparisons between different groups of data entities. Visualization of the embedding space allows the user to observe the comparisons directly [HG18, GHM21]. With contextualized information, comparisons can be made at various levels. As previously mentioned, pre-trained embeddings offer multiple layers of information, providing a space for user interactions in a hierarchical fashion. Conducting multi-level explorations is a common VA interaction technique, but with embeddings, guidances can be aggregated based on the similarities of words in the embedding space [JWC*23]. Further, embeddings can be calculated based on the same or different context of a word, providing interactions with different levels of granularity (Figure 6(a)) [SKB*22].

Compared to exploring the resulting set of data clusters and the embedding space, for other techniques, users are more active and participatory in human-in-the-loop language model optimization. For example, El-Assady et al. model semantic spaces based on a set of word embeddings and visualize concepts at different semantic abstraction levels for topic model refinement [EAKC*20]. Sperrle et al. provide active guidance in the topic modeling process. They map the embedding of keyword descriptors to visual cues in the interface [SSKEA21]. For data-driven techniques, user interaction with embeddings offers more support to specific applications such as peer reviewing [KOK*18] or general document labeling systems [RSBV21]. Here, the guidance from embeddings is beyond ways to explore the semantic space—it supports users in making decisions on a higher-level analytical task.

Handling uncertainties from NLP-related tasks is another motivation for several papers in the set. From a word disambiguation perspective, a task of uncertainty tackling could be providing an interactive ambiguity resolution technique for named entity recognition [SJB*17], readjusting semantic relation of concepts based on users' understanding for topic model refinement [EAKC*20], or supporting an information retrieval system to achieve high recall [DMdO19]. Adding embeddings in a VA pipeline not only represents input text data, but the projection of the embedding space also acts as a source of validation to the process of word disambiguation [DMdO19].

Neural Network Interpretation We observe less direct user involvement for general neural network visualization papers in the corpus. Like neural language models, generalized neural network-related VA tools in our corpus focus on model understanding and validation [BPP*16, SGR*20, FZCM20, CDHP21, RSL*22, PDD*22, LWBM22, SGL22]. Users may explore decision and feature spaces for counterfactual reasoning [SGL22], kernel actions [BPP*16], and neuron group interactions [PDD*22], or validate unfairness and errors in AI systems [CDHP21, RSL*22].

In many cases, embeddings may not function as a guidance generator, and users may not actively participate in a human-in-the-loop process to optimize a model. Instead, visualization of embeddings aims to give an overview of the underlying data distributions. Global structures are presented to aid model understanding and validation, rather than a details-on-demand visualization of local data subsets. Correspondingly, the predominant visual representations are 2-dimensional scatter plots projecting the embedding space to give an overview of the dataset. One particularly interesting ap-

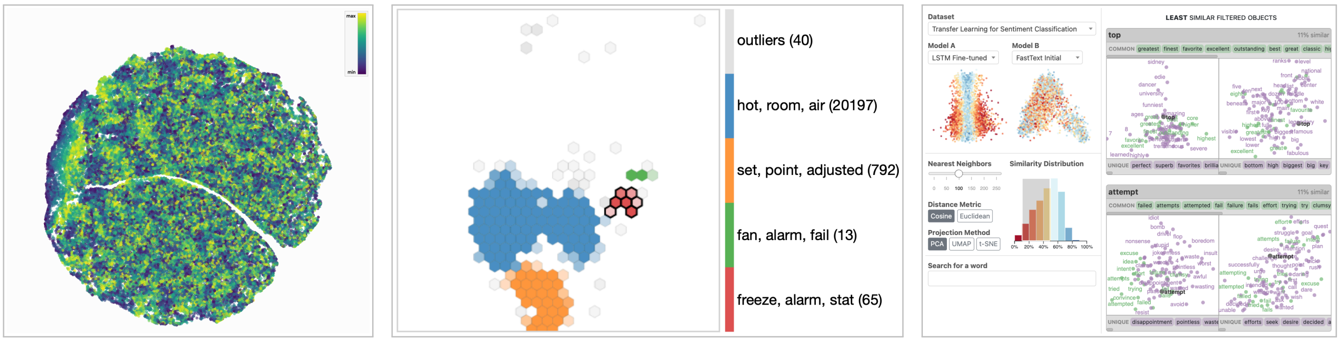

(a) (b) (c)

Figure 5: *Examples of 2D-projections of the embedding vectors. (a) Cytosplore by Höllt et al. [HPvU*16, vUHP*17] for visual analysis of biological data. (b) Visual analytics approach for machine maintenance data by Zhang et al. [ZFC*21] (image courtesy of Xiaoyu Zhang). (c) Embedding Comparator by Boggust et al. [BCS22] for interactive comparison of user-defined embedding models.*

proach in this set is from Liu et al. [LWBM22] (see Figure 7(a)). They present node-link diagrams to visualize node embeddings for graph neural networks. This technique captures both global topologies and features of the data in the visual representation. Numerous other VA-assisted ML (VIS4ML) [SZS*17] papers address similar motifs as embeddings.

As stated in many VIS4ML surveys, visualizing the hidden state representations of neural networks appears almost everywhere in explaining the model [CMJK20]. However, they may refer to it using different terms such as “latent space”, “latent representations”, “feature space”, “activation vectors”, and “spatialization”, among others [WZY20, FZCM20, RCPW21, LWBM22, PDD*22]. Each of the terms works under a slightly different context. In this study, we have elected to focus on the VA papers where authors explicitly mention “embedding” and exclude others that may cover a similar topic but use different terminology. It should be acknowledged that this approach may introduce bias into our selection criteria. However, we believe the insights derived from our corpus are valuable and worthwhile. The findings will be further discussed in Section 7.3.

5.1.2. Computing

In this category, we include papers that focus on computing-related problems and applications beyond ML. The reason for this design is two-fold: first, as mentioned above, we drew inspiration from ACM CCS and the categorization of application domains used by IEEE VIS / VGTC, with the latter separating ML-related applications from other areas of computing (e.g., databases, computer networks, or security). Secondly, given the topic of this STAR and the expectation for a large number of publications/techniques focusing on VIS4ML, etc., we preferred to keep this category separate in order to paint a more clear picture of the respective applications. Most of the papers in this set focus on mining graphs and/or sequential data. The embeddings may represent the nodes of a network [CZIM18, XXM19, XTYL20, PCZ*21, LTHL21, SDXR22, LWBM22] or sequences of events [LDL*20, XTYL20], ultimately aiding in identifying structural, temporal, and multivariate properties within groups of nodes [PCZ*21], communities [CZC*17], ensembles [FFST19, WJM*22] (see Figure 6(c)), etc.

Performance and Software Visualization For performance visualization in high-performance computing, a call stack tree can depict the context of function execution and retrieve anomalous execution behavior [XXM19]. Similarly, creating a visual analytics tool for software engineering data may involve understanding complex code dependencies. Constructing a hierarchical graph using file structure can assist in obtaining statistical information and evaluating the similarities between bad dependencies [LTHL21].

When we review the VA pipeline, in performance visualization, embeddings represent a variety of data structures, including file structures, source code dependencies, and other data ecosystems. The dataset can be heterogeneous. Computing embeddings is often at the beginning of a VA pipeline to unify input data into the same vector space, particularly for anomaly detection and risk inspection. In *PRIVEE* [BIVD22], embeddings are used to find and represent joinable datasets. Additional operations, such as weighting factors, can be incorporated into embedding vectors to convey additional customized, attribute-based information.

For those tools, the visual representation of embeddings is a supplementary support of navigation for the main panels. Even though graphs are embedded, the embeddings are represented visually by scatter plots rather than by node-link diagrams. Positions of the points convey semantic similarities of the embedded entities. Users can use the 2D projection of embeddings to filter out or select points that may appear anomalous compared to others.

Visual Search By embedding the data into a high-dimensional space, it becomes possible to identify patterns and construct interactive search query interfaces. For either topic-, graph-, or image-based queries, it can be challenging to create a search motif that is both intricate and generalizable. Furthermore, the corresponding interface needs to consider providing guidance, expansion hints, and feedback to user-defined queries. The commonality between *TopicSifter*, an interactive search space reduction technique [KCD*19], the interactive visual pattern search technique proposed by Song et al. [SDXR22], and *VISAtlas*, an imaged-based query system for visualization collections [YHZ22], is about finding ways of creating human-interpretable embedding vectors. Users can observe object alignments inside the embedding space. At the same time, those clusters guide users in refining their search.

A visual querying interface often consists of three key components: the original data, the queries, and the filtered results. Neural network-based embedding methods enable data from different modalities to be projected into a vector space [GWW19, YHZ22]. For many applications, measuring data similarities in the embedding space is the most essential part of querying, but there are multiple ways to interact with the space to obtain the final querying result. Filtering can be restricted to documents containing keywords, to keywords and their nearest neighbors, or loosened to include similar documents ranked by some distance measure in the embedding space, giving users multiple levels of constraints with thresholds to customize their search. Embeddings are often integrated into the computational part of a VA pipeline and may not be explicitly mapped to a visual representation [KCD*19, SDXR22]. In other cases where embeddings are explicitly represented, the interactive interface allows for another dimension of user customization. For example, in *VISAtlas* [YHZ22], the 2D scatter plot projection of the embedding space allows users to query similar data objects by selecting points on the panel. Regardless of whether embeddings are visually represented or not, users can have multiple ways to filter and query items by changing the distance metrics, adding constraints to the aggregation and filtering process, and providing a target set with additional threshold values.

There are various novel visual encodings available in this category, especially for layout refinement in graph drawing [CZIM18, XTL*21, JCS*21, PCZ*21, TCS*22]. Our visual browser includes a filter button for **computing** in the application domains, allowing readers who are interested in comparing those tools directly, and potentially they can be a source of inspiration.

5.1.3. Life Sciences and Medicine

Mainly papers in biology and medical natural language processing (MedNLP) are included in this set.

Biological Data Visualization For those biology visualization techniques, data are collected from massive individual cells [HPvU*18] or epigenetic modifications within the human genome sequence [LPH*20]. The size of the data is large and high-dimensional. For instance, single-cell data could be gigabytes to terabytes [LGY*20], and the human genome is a sequence of roughly 3.3 billion chemical units [LPH*20]. Researchers seek to summarize those large-scale data into an exploratory lower-dimensional space and find patterns that shed light on some biological phenomena. After applying embedding techniques, the learned representations can be utilized to generate clusters at different granularity levels. Valuable insights can be obtained by visualizing hierarchies of clusters [HPvU*16,HPvU*18] (Figure 5(a)) or embedded representations of omics data sequences [AHH22].

Unlike many other domains where visual representations of embeddings are often supplementary views, embedding representation usually plays the central role in this context. Due to the large scale of biological datasets, computational complexity is a major concern when integrating these methods into a VA pipeline. Clustering derived from embeddings plays an essential part in this set of biological data visualization, as different aggregation levels of clusters provide different levels of summary view for the large dataset. Here, inspecting the points, clusters, and the additional associated information is enough to develop an effective VA system for a research question. Typical tasks of exploring the embedding space include examining the feature characteristics of each cluster, discovering inter-cluster relationships, and associating observations from the feature space to other reference data. However, variations of those tasks are common in other domains as well. As such, we link them with high-level visual analytical tasks in our categorization elaborated in Section 5.6.

Visualization for Healthcare MedNLP is a subfield of NLP that often uses language models to extract representations and contextualization from medical records, medical claims [CEBV22], treatment concepts [JSR*19], and patient information [KGM*22]. The included VA+embeddings papers can also be considered a subfield within the healthcare domain. Compared to those NLP papers previously mentioned in the subsection of AI/ML domain, we can verify that these MedNLP papers focus primarily on leveraging users' domain expertise as inputs for a VA system. Focus has shifted from exploring linguistic dimensions to constructing a patient network from raw data [JCS*21, KGM*22] or modeling patients' histories for clinical prediction [KCK*19, LYY*20].

Based on electronic medical records (EMR) or electronic health records (EHR), it is common NLP practice to extract features as low-level building blocks and identify topics as high-level syntheses of documents sharing similar characteristics [JSR*19]. Further, documents and topics guide the exploration of features and embedding dimensions. However, for medical data, it is even more critical to ensure the interpretability of a system. A clinician with sufficient domain knowledge can explain the reasons behind their diagnosis to a patient. A VA tool is expected to incorporate a similar level of interpretability not just for patients but also for clinical professionals [KCK*19]. Electronic health and medical records are often large-scale, sensitive, and irregular in terms of timelines. Many papers not only work towards summarizing them accurately, but also aim to produce a visual analytics system that clearly illustrates how and why a particular prediction is made [JSR*19, KCK*19, LYY*20]. If the produced representations are similar, it is often necessary to communicate in which aspects they may be different [GFL*20, JCS*21, KGM*22, CEBV22]. The emphasis on high recall and system credibility in meeting user expectations has led to changes in research papers' framing and evaluation methods. We elaborate on this point across multiple domains in Section 5.2.

Oubenali et al. conducted a scoping review on visualizing word embedding explicitly for medical concepts and identified 7 papers within this scope [OMF*22]. Their work confirms our analysis that visualization is used to explore embedding results for MedNLP applications. Despite many papers highlighting the importance of trustworthiness in this domain, there is still work to be done in evaluating the embedding-related systems with respect to both computational performance and the effectiveness of the visual representations. In Section 7.1.4, we will discuss the open challenge of conducting an evaluation of VA systems under different contexts.

5.1.4. Humanities, Social Sciences, Education

This set contains VA tools on urban science [XTYL18, LKJ*20, MZAD*20, MHL*20, BZQ*21, RMH*22, GZRP*22, SNP*22] and for the analysis of various sources of data, including social media [BEF17, XO21, WSP*21, AAM*21, AYL*22], news/creative writing [PS21, HGE22], literature/digital humanities [NKWW22, MWJ22], and human behaviors/gestures [WLHO19, ZWW*22].

Incorporating and representing temporal information is one of the primary motivations for using embeddings within this set. Even though there is a large design space to visualize time series from different narratives [BLB*17], the irregular and multivariate nature of many datasets in this category makes it challenging to effectively present temporal data without abstraction and vectorized representations. Furthermore, a common goal is capturing similarities between phrases or sequences. Even though dynamic time wrapping (DTW) [Mül07] is an alternative method, embeddings generated from models such as autoencoders [KW19] and word2vec [MSC*13] are useful as they aim to capture semantic similarities by transforming the data into features. While there are more granularities between embeddings and other alternative methods when it comes to integrating them into a VA pipeline under a similar context, neural embedding models are often considered a more general approach for both representation and comparison purposes [GZRP*22]. The embedded space can be visualized using 2D scatter plots, enabling users to track the embedded vectors to the data entities, along with a temporal view and sometimes a spatial view. Further discussion of the computational methods being used across different domains can be found in Section 5.4.

Social Media Visual Analytics Visualizing social media data is an interdisciplinary field, as it involves applying language models for posts, performing network analysis, and applying text and multivariate visualization techniques. For those VA+embeddings papers, they do share a common motivation in developing the system—social media is populated with false and harmful information, and their tools aim to assist in the decision-making process. For example, *Recast* [WSP*21] uses interactive visualization to help increase the interpretability of toxicity detection models. Przybyła et al. [PS21] propose sets of visualization tools to explain the news credibility assessment process interactively. Other papers do not directly focus on misinformation, but they do mention that understanding social network topics and their transitions plays an important role in the decision-making of government [XO21, AAM*21] and individual analysts [BEF17]. Embeddings are often produced during the neural network classification process, and visualizing them allows users to find data sharing similar classification/prediction results, such as harmful information. On the other hand, users can rank the distance between a source object and the rest to produce a list of items as potential alternatives. Correspondingly, ranked lists, tables, and word clouds are displayed in the interface.

Even though the high-level objectives sound similar, this set of papers presents a variety of customized visualization techniques. *Recast* [WSP*21] and the work from Przybyła and Soto [PS21] have simple yet intuitive visualization consisting of one scatter plot or a few bar charts to help with language model explanation. On the contrary, the VA systems proposed by Xiao et al. [XO21], Badam et al. [BEF17], and Andreadis et al. [AAM*21] embed multiple customized and juxtaposed views with close and distance reading to represent topic relationships, temporal trends, geo-information, and user interactions. The complexity of visual interfaces can vary depending on their research question formulation, but embedding-related views generally rely on 2D scatter plots as the predominant visual representation. However, lines/links are often used as an augmentation approach to convey relationships between data entities, features, and other attributes. The word "alignment" is frequently used to describe the motivation for integrating embeddings [MWJ22, XWX*22, AYL*22], and computing clusters under different time snapshots or other constraints can inform changes or connections. To facilitate the interpretation of clustering results, parallel coordinate plots [XWX*22, CGH*22, AYL*22], line charts [GZRP*22], and node-link diagrams [XO21, CGH*22] are introduced as ways for users to find alignment for individual features, multidimensional attribute values, or within-cluster and between-cluster trends and patterns.

Urban Visual Analytics In terms of visualizing urban data, designing visual representations and interfaces can help better understand the characteristics of a city, such as human mobility patterns, ultimately enhancing the lives of citizens and promoting sustainable development [RMH*22]. However, the challenges of analyzing urban data are highly dependent on specific techniques, since data may come from multiple sources, such as maps, traffic flows, human movement trajectories, loop sensors, acoustic sensors, social applications, etc. Although the generalized workflow may consist of data retrieval, parsing, analysis, evaluation, and visualization [Hu18, MZAD*20], the specifics are highly dependent on the desired use case and context of research questions. For example, the design considerations for a real-time monitoring system would be different from an offline analysis tool. As a result, there are multiple considerations for embeddings and diverse ways that embeddings are integrated into a VA pipeline.

In the following paragraphs, we provide a brief overview of three key areas for those urban visual analytics papers in our corpus: crime pattern analysis (visualization for public safety), traffic congestion monitoring (visualization for traffic flow), and geo-text data visualization. We also discuss the different roles embeddings play in those three areas, but it is not surprising that embeddings mainly function as a bridge that links original irregular and heterogeneous data entities to specific spatial areas for users to explore on the map, and a common challenge is integrating those additional supports, e.g., clustering, classification, or prediction results, onto the map.

Visualization for Public Safety Crime data, which is crucial to gather for enhancing public safety, is often sparse and spreading in large spatial areas [GZRP*22]. Identifying crime hotspots and uncovering patterns requires gathering a combination of temporal sequences, geographic locations, socio-demographic attributes, and other descriptive categorical values. *Crime Prevention Through Environmental Design* (CPTED) [CSH05] often focuses on finding statistical correlations between elements of the environment and crimes. *CriPAV* [GZRP*22], on the other hand, uses spatial discretization to convert streets and streets' intersections into networks, and employs autoencoder to embed crime time series in a Cartesian space. Crime hotspots are detected by applying a clustering algorithm for the resulting feature vectors. The proximity information in the embedding space is one important factor that influences the choice of users to select a point/place as anchor points to further link with Google Street View photos. Neural network-based embeddings are generated and projected into a 2D panel to classify events concerning public safety, but at other times they are the byproduct of predicting the ongoing temporal sequences, especially in the area of visualizing traffic flows.

Visualization for Traffic Flow Aside from dividing urban areas through spatial discretization that divides a physical continuous space into discrete regions, for analyzing traffic flow, it is important to consider both inter-area directional interactions and local attribute information [ZDL21]. With spatial-temporal data as inputs, traffic flow and congestion are often derived from transit mobility structures such as bus stops. Graph embedding is one way to represent those geospatial map-like data after discretization. Lee et al. [LKJ*20] employ DeepWalk [PARS14] to embed a road network to extract latent similarity features. As part of their approach, the embedded feature matrix with speed and rush hour matrices will be fed into a Long Short-Term Memory (LSTM) [HS97] model for the classification and forecasting of traffic congestion. Zhang et al. [ZDL21] employ attributed graph embedding for transit trips. They combine mobility pattern graphs and attribute similarity matrices into an autoencoder model [VLL*10] to detect mobility communities. To understand the temporal-spatial correlation, many works focus on either embedding time series or discretized maps as graphs, such that correlations can be represented by map coloring, or observed by placing marks on a map.

Geo-text Data Visualization Besides crime hotspot detection and traffic flow analysis, other geographical visualization in this set focuses on spatial-temporal topics coming from geo-tagged product reviews [XTYL18] and social media posts [MZAD*20]. The corresponding VA tools embed only the textual data, and tend to focus more on integrating the resulting word embeddings that reveal linguistical dimensions with a map in a 2D interface.

Despite the limited number of papers in this set, we observe a few intriguing connections. For the included urban/geographical visualization techniques, the VA tools focusing on developing complex embedding methods for representing data tend to provide multiple linked panels to explain the embedding projection results. Meanwhile, papers that deal with simpler embedding techniques, seem to prioritize building an intuitive and effective visual interface with overlay for users to explore. In the paper *GTMapLens* [MZAD*20], Ma et al. propose a lens-based visual interaction technique that provides word recommendations to users. When a user explores a geographical region and inputs a keyword of interest, the content of the lens will automatically trace and highlight closely related words. *GTMapLens* uses word embedding to generate word recommendations. It would be interesting to investigate the possibility of incorporating other types of data, such as time series, into the embedding process. At the same time, *CriPAV* [GZRP*22] embeds crime time series for users to detect crime hotspots. They provide a location view with dots to show the detected crime hotspots along with multiple linked panels displaying temporal trends. It would be of interest to enhance user interaction by allowing them to input a keyword or timeframe of interest, and having the system trace and suggest related concepts. This is just a simple and brief hypothetical scenario. There are certainly more opportunities for future directions when connecting VA+embeddings tools across different domains. We will further discuss the trends and opportunities in Section 7.

5.1.5. Business, Management, Governance, Law

Research on visualization for business analytics and management tools focuses on developing interactive interfaces for various business and management data. There are several topics covered in this set of papers, including market data analysis [CKC19, CKC20, HKD*21, JCSM22, CKN22, CGH*22] (see Figure 7), collaborative and multi-user analysis [XBL*18], visualization of meeting content [CBS*19], and machine maintenance data [ZFC*21] (Figure 5(b)).

Market Data Visualization For businesses, a market data visual analytics tool can inform customer satisfaction with their products. It can be used for strategic recommendations in marketing and decision-making purposes [CKC19, CKC20, CKN22], as well as identify marketing targets [CGH*22]. Meanwhile, VA techniques can help customers explore product-related topics [HKD*21] (Figure 7(c)) and perform serendipitous discoveries [JCSM22]. A popular source that produces market and product data is social media platforms, containing large-scale text and network data. Although most of the papers in this set are not directly related to natural language processing, they express the need to use state-of-the-art NLP methods to perform sentiment, opinion, and aspect analysis [CKC19, CKC20]. In such cases, word/sentence/document embeddings are calculated as part of the computational pipeline to output a correlation/similarity rating for topics or sentiments, and there is either no explicit visual representation or, for instance, a bar chart or some percentage numbers as extra information.

Visual Interfaces for Collaboration Designing a visualization tool for collaboration can be synchronous or asynchronous. For synchronous group support systems, data is shared, analyzed, and interacted with by multiple persons in real time. Meanwhile, asynchronous VA tools enable sharing and reviewing prior workers' outputs. Effectively summarizing and connecting findings are important for both types. Via embeddings, multi-user outputs can be represented to calculate "relatedness", such that similar contents can be put together in the interface [ZFC*21].

Depending on the usage scenario, the collaboration object to be embedded may vary. For example, *TalkTraces* [ZFC*21] provides a visualization tool for displaying real-time meeting contents. Word embeddings are used to determine how discussions relate to one another. *Chart Constellations* [XBL*18], on the other hand, supports a single analyst to review and analyze data visualizations from prior collaborators. They also incorporate word embeddings of tagged keywords in the meta-label, along with chart encoding and dimensional intersection information to aggregate pairwise distances such that similar charts can be clustered. Even though the visual interface is still largely encoded as 2D scatter plots, each point now rep-

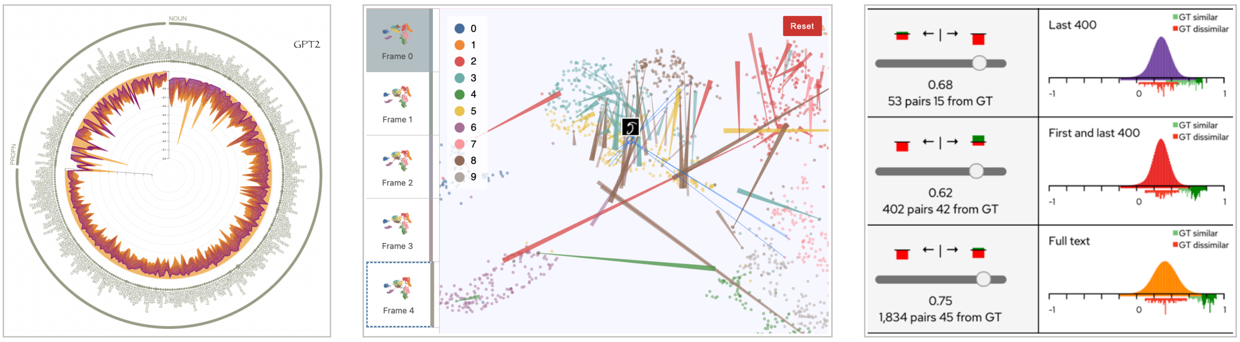

(a) (b) (c)

Figure 6: *Examples of custom visualizations of embedding data. (a) LMFingerprints by Sevastjanova et al. [SKB*22] for visual explanation and comparison of embedding spaces. (b) Emblaze by Sivaraman et al. [SWP22] for animated and interactive comparison of embedding spaces with Star Trail augmentation. (c) EEVO by Witschard et al. [WJM*22] for interactive optimization of embedding ensembles.*

resents a visualization chart. An ensemble calculation of different distances with user-customized weights determines their positions.

5.1.6. Physical Sciences, Engineering, Mathematics

The papers in this set focus on exploring model collections [AKZM14], analyzing air quality evolution by embedding pollutant event sequences into latent stages [QLL*22], extracting contents of mineral exploration reports by embedding keywords and visualizing semantically similar words [WMW*22], and designing new lubricants by calculating importance and correlation values from molecular embeddings [MNS*23].

Visual Browsing of Multimedia Collections Shape browsing and navigation refer to a process of exploring collections of 3D shapes, images, labels, and other multimodal representations through an interface. It falls under the intersection between computer graphics and human-computer interaction. More often than not, many papers on this topic have a very simple interface with minimal visual representations, often a 2D projection plot, but it is still relevant to our survey. There exist two main challenges. First, given a high-dimensional search space of images and shapes, it is not efficient to generate a global manifold that preserves similarity relations among all objects for a user to navigate [KFLCO13]. Second, given the low-dimensional representations, users need to effectively navigate through the space and perform desired tasks in an intuitive way [AKZM14].

There are many papers aiming to address the first challenge. Since users essentially are browsing through a small subset of objects, one approach is to generate a local manifold that only captures the region currently observed by the user [KFLCO13]. Another method, which is more of a global approach, is to apply a neural network to extract feature vectors for all the objects, and then employ a DR method to project features into a 2D space with a 2D projection plot for exploration [LSE21]. For the second challenge, due to our restricted and limited search (as described in Section 4), fewer papers aim from a visual analytics and interaction perspective. One particularly interesting example, *ShapeSynth* [AKZM14] not only provides an embedding technique for part-aware shape descriptors, but also presents an interface for exploring the embedding hierarchically at various levels. The overall application goal of this approach is to enable non-expert modelers to explore possible synthesizable shapes for their model. We expected other related papers to narrate from downstream domain applications' perspective, and aim at different audience groups within the physical sciences and engineering domain. However, we were not able to identify further papers that fit the scope of this survey in this regard so far.

5.1.7. Sports and Entertainment

There exist a few papers that make our corpus a little more diverse. We include them here. The topics of those papers involve game visual analytics [XWX*22], sports visualization [WWC*21], and an interactive system for art generation [FCH*22]. Those papers use different embedding techniques and different visual representations of the resulting embedding. For *RoleSeer* [XWX*22], a visual analytics tool to inform social role changes during gameplay, dynamic network embedding is used to represent the social identity of players. Users can observe role transition by comparing clusters at different timestamps. For *Tac-Miner* [WWC*21], a visual table tennis tactic mining system, table tennis strokes are embedded via word2vec as a basis for further visual analytic tasks. For *iPoet* [FCH*22], a multimodal system to allow users to compose poems for paintings, embeddings play a role in labeling and visually showing the sentiments of poems, and they are generated via gated recurrent units. Even though the contents between those VA tools vary a lot, they still use the same building blocks as other domains: using word2vec variants to embed domain-specific data, such as the strokes from racket sports, and calculating graph embeddings to represent social relationships, similar to what has been discussed in Section 5.1.4.

5.1.8. Domain-Agnostic

We define a VA approach in our corpus as "domain-agnostic" if the system focuses on embedding space visualizations. In such cases, the source of the embedding space may be user-defined, unspecified, or independent of domain-specific tasks. There are two possible scenarios for papers to be included in this set. (1) Papers may be framed in a specific context, but the input data for embedding are generalizable. For instance, *Emblaze* [SWP22] is designed for ML representations, but the emphasis remains on embedding space exploration, and they do not explicitly mention the constraint of input data (see Figure 6(b)). Another similar example is from Renoust et al. [RRM*21]. While their VA system is largely focused on information retrieval and querying, the embedded data can be generalized to support multiple formats and types. (2) For systems that only deal with one data type, such as text, their system can work as a module for tasks across various domains. For instance, Witschard et al. [WJM*22] provide an interactive optimization system focusing on text embedding-related metrics (Figure 6(c)). Their system is applicable to fields beyond the scope of text analysis, though.

Visualization for Embedding Interpretation Embeddings are ubiquitous in machine learning. Beyond designing a panel for embedding analysis in a larger system of ML model interpretation as mentioned in Section 5.1.1, there exist many visual analytics papers dedicated to interpreting user-defined embeddings [STN*16] or exploring user-defined multi-dimensional data via embeddings [XTL*21]. Smilkov et al. [STN*16] summarized three high-level tasks to facilitate the interpretation of embeddings: (1) exploring local neighborhoods, (2) viewing global geometry and finding clusters, and (3) finding semantically meaningful directions for a certain concept set. A handful of papers in our corpus from various domains reflect at least one of the themes.

VA systems can be used for comparing two or more embedding spaces in addition to exploring one embedding space. For those types of techniques included in our corpus, it seems that some of them choose to focus primarily on one of those three high-level tasks. Extending task (1) to compare local neighborhoods, *embComp* [HKMG22] provides a neighborhood overlapping view displaying the number of similar neighbors in the two embeddings, along with other metrics reflecting properties such as spread and density of neighbors. Extending task (2) to compare the global geometry of embedding spaces, *EmbeddingVis* [LNH*18] provides multiple panels in a row to compare the global structure of different embedding results. Task (3) of identifying semantically meaningful directions may be of interest to computational linguistics researchers. Especially for word embeddings, subtracting the vector representations of two related words and adding the result to a third word would help researchers detect potential biases in those vector relations [BCjC19]. A system from Liu et al. [LBT*18] aims to address this task primarily. The proposed tool supports visual exploration and reviews the semantic relationships of syntactic analogies in an embedding space. Even though there are no side-by-side comparisons in the interface, users can switch between word embedding methods. *Embedding Projector* [STN*16], which specifically works on interpreting embeddings interactively, supports users in performing all three tasks. In Section 5.6, we link those three tasks with the high-level visual analytic tasks in our categorization schema suitable for all domains.

As the complexity of VA applications increases when it comes to comparing multiple embedding spaces, there is tension between adding more features to the visual interface versus making the interface simpler and more intuitive to enhance visual literacy [Rus16]. As a result, it is crucial to find the right balance for the targeted audiences. Visual interfaces need to provide them with enough perspective to facilitate their tasks of interest and enough evidence to make the sophisticated metrics and visual representation trustworthy. We will discuss these trade-offs in detail in Section 7.2 *Visual Analytics for Embeddings*.

5.2. Target User

Visualization tools assist users in performing tasks such as discovering trends or identifying outliers. It is equally important to consider the profiles of target users, which in turn have an important effect on the chosen design. However, it was unexpected to discover that only 32 out of the 122 papers provide **explicit target user description**. Out of these, an even smaller portion explicitly stated the assumed level of user expertise [WWC*21, FCH*22, SWP22, AHH22]. While these numbers clearly should prompt some self-scrutiny within the research community, it must also be said that many of the publications contain a lot of implicit information, which makes it possible to deduce a profile of the targeted user. Hence, the situation is not as bad as it first might seem.

Two common ways for papers to mention the scope of targeted users are expert interviews before development and post-development evaluations, including usage scenarios, user studies, and quantitative measures. While evaluation concerns are discussed in Section 5.9, there are two types of assumptions for defining target users at any stage of development beyond formulating implications and design goals. First, all users share a set of common knowledge, or at least they need to meet certain criteria. Similar to many other VA papers, a common reference to indicate such shared knowledge is to refer to the users as "researchers" [KCK*19, CEBV22, RSL*22, SCR*23], "domain experts/scientists/specialists" [LBT*18, AHH22, MNS*23, WHC*23], "expert readers" [MWJ22, JCSM22], and "practitioners" in a certain domain [MHL*20, XWX*22, RSL*22, SWP22]. Even though those terms provide a narrower scope than simply "users" in order to refer to the target groups of a VA tool, still, the explicit prerequisites to use a domain-specific tool are hidden behind the scenes.

Besides commonalities and user study guidelines for developing every VA tool, we observe different levels of requirements for using embedding-related panels. Obviously, for tools that aim at exploring embedding spaces, users are expected to have experience in working with embedding models. While other tools incorporate an explicit visual representation of embeddings as a supplementary panel, several authors mention the following criteria when deciding on user study participants. Even though the target users may be from a specific domain such as medicine or just game players, they are implicitly required to have "basic concepts of machine learning", use "software for patient data analysis", be "computer science major students", or students who understand "statistics" and "intrinsic dimensionality" [XWX*22, CEBV22, MM23, WHC*23].

As a result, we would like to highlight the user-centric eXplainable Artificial Intelligence (XAI) considerations proposed by Wang et al. [WHC*23]. Even though not all embedding methods we surveyed are AI-based, according to the well-defined considerations from Wang et al., we believe it is valuable to specify users' research

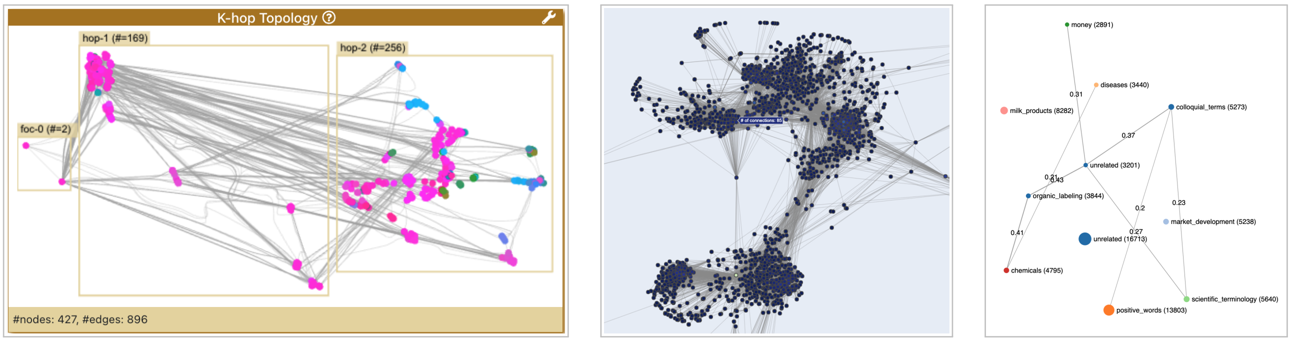

(a) (b) (c)

Figure 7: *Examples of node-link diagrams for trees and graphs/networks, mainly or partly derived using embedding data. (a) CorGIE by Liu et al. [LWBM22] for visualizing graph neural networks. (b) BiaScope by Rissaki et al. [RSL*22] for interactive investigation of unfairness for graph embeddings. (c) SocialVisTUM by Hagerer et al. [HKD*21] for exploration of relationships between topic embeddings.*

field, AI (in this case embedding-related) expertise, and their role in using VA+embeddings systems when writing a paper.

Creating embedding-related panels may require users to have a certain level of knowledge to use the tool effectively. However, from another perspective, embeddings open up opportunities for creating personal visualization [HTAA*15] that enables data analysis in a personal context. *Metaphorical Visualization* by Tkachev et al. [TCS*22] proposes an inspiring VA+embeddings approach that uses embeddings to create metaphors. They link one data entity to another via distance-based mapping derived from ML embedding spaces along with other mapping algorithms. As an example, authors can be explored metaphorically by mapping them into English nouns or even cat images. Although users may have different backgrounds and goals, the created visualization aims to encourage all of them to experiment with it. As discussed in Sections 5.1.4 and 5.1.5, several other approaches also use embeddings to support serendipitous discoveries. As noted in the evaluation section of the work from Jasim et al. [JCSM22], users with explicit knowledge about a data item also gain implicit knowledge about all other data entities that are semantically similar. For such cases, embeddings help to mitigate the knowledge gap by indicating semantic similarities within and across different datasets, allowing users to infer unknown concepts based on their existing background. At the same time, such approaches ensure less redundancy and a more balanced exploration of data with different attributes.

We want to summarize this subsection by reiterating four observations. First, many papers in our corpus have not explicitly mentioned target users and usage contexts for their proposed VA tools. Second, the scope and background knowledge of target users that authors frame in the introduction of their paper is often different than what they illustrate as criteria for finding suitable candidates for user studies. Thirdly, introducing embedding-related panels adds further background assumptions to target audiences, and papers rarely mention such assumptions. Lastly, embeddings open a broad space for personal visualization [HTAA*15], as one of the main motivations for using embeddings is to make distance-based recommendations to potential exploration targets. Incorporating embeddings to make guidance can leverage users' knowledge from one domain to another if the concepts are represented by a common embedding space [TCS*22].

5.3. Input Data Type

One important factor in assessing the use and spread of embedding technologies within the visualization community is to keep track of the type of data which is being embedded. The classification scheme used for this survey (together with the number of publications within each class) is: **textual** (70), **numerical/tabular** (25), **graph/network** (22), **mixed** (17), **image/video/audio** (17), and **data-agnostic** (6).

Since NLP is one of the domains where embedding technologies were first adopted and their use and development are still highly successful, it is no surprise that **textual** data input is used within almost half of the publications in our data set. Prominent examples of use from this domain are semantic analysis [HG18, MWZ19, GHM21, LWZ*23], document search and retrieval [CWDH09, KOK*18, RSBV21], and analysis of social media posts and/or customer reviews [BEF17, WSP*21, CKN22, JCSM22].

The second largest category, **numerical/tabular**, accounts for nearly one-quarter of the publications. Here, it is important to note that this category generally includes much more intricate examples than just collecting values over several numerical features. For instance, there are examples of analysis of sequential data (time series, mobility patterns, etc.) [LPH*20, BZQ*21, ZJQH22] as well as of applications that load embedding vectors that have been calculated and published by other parties [HKMG22, EHA*22]. **Graph/network** data also accounts for an important contribution. This field has seen substantial growth after some initial word-embedding technologies were modified to handle this data type. Common usage scenarios are analysis and comparison of graph topology [CZC*17, PCZ*21, SDXR22], and (just as for word embeddings) comparing the embedding spaces of different algorithms can give important insights to the different algorithms as well as to the underlying network [LNH*18, CZG*22]. Furthermore, TorusTraffic$^{\text{ND}}$ [CZIM18] provides an interesting example of embedding network topology onto a Hilbert curve rather than using a numerical vector.

As indicated by the name, the **mixed** category includes examples of applications that use more than one input data type. A common scenario is to combine textual data with time and/or geospatial data [XTYL18, GFL*20, LYY*20, AAM*21], and another approach is to combine network topology with node attribute data [LHZ*22]. As for the **image/video/audio** data type category, the most prominent example is image data, which is used both on its own (e.g., in applications using image/shape similarity calculations) [KFLCO13, AKZM14, YHZ22, HHS*23] as well as in combination with other data types [PdSP*22, GGW22]. An exception to the image-based focus of this category can be found in [RMH*22], which embeds audio data for analyzing urban noises.

Finally, we have the **data-agnostic** category, where the input data is typically treated as generic numerical vectors, and little, or no, assumptions are made on the underlying data. Here, we find applications for more generic exploration and comparison of different embedding spaces without taking full advantage of knowledge from the underlying data domain [STN*16, BCS22, SWP22].

5.4. Computational Method

We classify the computational approach into three main methods, which can of course be used in combination within the scope of a single paper. The most common, with 88 occurrences, is **neural network**, where different types of ML/DL models are used to obtain the embeddings [ZJQH22, RSL*22, SDXR22]. Even though there are many different variants, some of which are highly customized for the specific tasks, our analysis shows that using the BERT [DCLT19] model or word2vec [MSC*13] (or some of the other versions inspired/derived from it) is a common choice among the surveyed papers. This is yet another indication (compare to 5.3) that NLP continues to be a key area for using and developing embedding technologies, and that advances regarding textual data sometimes also can be transferable to other data types, such as the case of word2vec (used for word embedding) whose algorithm has been modified to obtain variants for embedding graph/network data, such as *stack2vec* [XXM19] for representing call stack trees, *struc2vec* [XWX*22] for learning the structural identity of nodes, and *metapath2vec* [DSS*23] for capturing cross entity interaction features. To explore all the survey entries using a variant of the word2vec techniques, readers can enter "2vec" in the search bar of the visual browser accompanying this survey article (assuming that the respective titles are available as part of free-text details).

The second most common method (with 33 occurrences) is **matrix analysis**, which relies on matrix decomposition methods such as spectral analysis/eigenvalue decomposition to calculate the embeddings [KCD*19, FFST19, PCZ*21]. Within this category, the choice of using the precalculated GloVe word embeddings [PSM14] is fairly common. The third category (with 16 occurrences) is **statistical analysis**, such as term frequency–inverse document frequency (TF-IDF) [KFLCO13, HMW*15, QLL*22]. Finally, the smallest category is the **computational method-agnostic** category, which accounts for 12 occurrences, and just as for the data-agnostic cases of Section 5.3, these applications do not make specific assumptions on the embeddings, but rather treat them as generic numerical vectors that are sometimes directly from user input [STN*16, HPvU*18, LJLH19]. From the overall distribution, we find it reasonable to conclude that the most prominent development of embedding technologies currently takes place within the ML/DL domain, as those neural representation learning methods can learn semantic, geometric, and contextual information. However, it is also clear that some of the "older" technologies still have their area of successful use.

After obtaining the embeddings, applying clustering algorithms and computing distance-based similarities for two vectors are often the following steps. While some papers focus on distances based on the embedding space, such as cosine similarities and calculating Euclidian distances [HKD*21, JCSM22], others may consider additional attribute values [GHM21] or graph centrality measures such as PageRank and betweenness [XWX*22] in combination with embedding proximity to make a final similarity ranking for data entities. Those distances and rankings can not only indicate semantically similar items but also suggest the fairness of an embedding approach. An example of such an application is *BiaScope* from Rissaki et al. [RSL*22], which visualizes fairness scores in graph embeddings for both individual nodes and groups (see Figure 7(b)). When using embeddings to provide recommendations, ensuring equal representation of all population groups is crucial. This can be achieved by measuring the proportion of embedding-based recommendations belonging to a specific population, which allows detecting if a sub-population is recommended disproportionately more [RSL*22].

5.5. Embedding Vector Dimensionality

The dimensionality of the embedding vectors affects both the possibility of capturing and representing the underlying data and the computational load (i.e., a vector of higher dimension may be able to capture the data in a better way, but on the other hand, calculations will be slower compared to using vectors of lower dimensions.) Since we expected word/text embedding to be a prominent topic within our corpus, and since many word/text embedding algorithms typically produce vectors of dimensionality within the range [50–500] dimensions, we chose this as the expected base bin of our classification and then added the side bins for shorter/longer vectors as well. From this, we obtain the following classification result: **under 50 dimensions** (15 entries) [STN*16, BCS22, PDD*22], **between 50–500 dimensions** (31) [CZC*17, KW18, XTYL20], **above 500 dimensions** (14) [KCK*19, WJM*22, RMH*22], and **output dimensions N/A** (70) [BEF17, MZAD*20, NKWW22].

As seen from the results above, many publications do not specify the dimensionality of the used embedding vectors. Even though authors often argue for their chosen dimensionality as "the best that preserves the semantics in the embedding space", there is often no proper justification provided. The choice seems to be based on empirical experiences and fine-tuning without a clear definition of under what metrics the results would consider "best". Furthermore, only a few applications have seen the need to use embedding vectors with more than 500 dimensions (which is mainly due to the fact that the BERT word embedding algorithm and the Universal Sentence Encoder typically output vectors of 768 and 512 dimensions, respectively). In general, the low attention that the discussion of dimensionality gets within the surveyed papers leads us to conclude that a main takeaway from this section is that vectors of up to 500 dimensions will most likely be sufficient and adequate (both in terms of data representation capabilities and computational load) for most applications, regardless of domain and underlying data type.

5.6. Visual Analytic Task

Regarding the classification of visual analytic tasks, we would like to acknowledge the prior work that laid a solid theoretical foundation within the fields of InfoVis and VA, including the work by Amar et al. [AS05, AES05], Brehmer and Munzner [BM13], and Schulz et al. [SNHS13], among others. For this survey in particular, we have been inspired by the VIS4ML ontology by Sacha et al. [SKKC19], but we have also seen the need for customization to fit our needs. The rationale is that we specifically want to target the reasons why embeddings are being used within the visualization—and this will not necessarily coincide with the (possibly) larger task that the whole application was developed. With the intent to provide answers to the question "What for/how are the embeddings used within this visual analytic application?", we have chosen to use the following categories: **model results representation** (63), **interactive exploration** (53), **comparison/selection** (52), **model results explanation** (23), **model construction** (14), and **model debugging/quality/bias control** (8).

Since embedding technologies have been designed to transform (possibly unstructured and complex) data and represent it as vectors, it should come as no surprise that the **model results representation** is the most common category. There are cases when this is the main task [BNL*18, BN21, PDD*22], but a common scenario is first to represent the underlying data and then use the resulting embeddings for data exploration [SJB*17, DMdO19, SGL22]. The same is in turn true for the category **interactive exploration**, which can be a main task [XTL*21, XTL*21, LTHL21] but also commonly be combined with comparison tasks [BMS17, BMS17, LJLH19, LZ23] (and with data representation tasks, as we have already seen). Furthermore, the comparison task may of course also be the only task [KW18, XTYL20, CEBV22]. With the current research focus on AI, it is logical to see that there are many cases when embeddings are used for model result explanation scenarios, both as a single task [BPP*16, GGW22, ZZL*22], but more commonly in combination with other tasks [STN*16, TWB*20, SSJ*22, SCR*23]. The content of the category **model construction** shows that embeddings can be used for constructing high-performing models [PKL*18, LKJ*20, WJM*22], and they may also be used for different model debugging scenarios [LXW*21, CDHP21, RSL*22].

In Section 5.1.8 **domain-agnostic**, we identified several representative papers that expand the three primary objectives when exploring the embedding space: (1) viewing local neighborhoods, (2) finding global geometries, and (3) exploring semantic meaningful directions for concept sets [STN*16]. Linking back to those tasks for domain-agnostic techniques, we would see that the categories we identified in this section can be constructed by a combination of those three fundamental tasks for visualizing embedding spaces.

The task **interactive exploration** can be applicable to both local and global embedding spaces. Corresponding to objective (1), some techniques in fields like *visualization for healthcare* (Sect. 5.1.3) and *performance and software visualization* (Sect. 5.1.2) choose to add visual representations to focus on viewing local neighborhoods as detail-on-demand or for anomaly detection. Corresponding to objective (2), for many other techniques in fields like *biological data visualization* (Sect. 5.1.3) and *urban visual analytics* (Sect. 5.1.4), the embedding views provide users a global overview of data distributions. Corresponding to objective (3), especially for fields such as *social media visual analytics* (Sect. 5.1.5), embeddings are visualized to reveal correlations of attribute features and user groups.

Moreover, comparing objectives (1) and (2) leads to the high-level task of **comparison/selection**, which commonly appear for *domain-agnostic* VA tools (Sect. 5.1.8) and *neural network interpretation* (Sect. 5.1.1) such as understanding contrastive neural networks [FZCM20]. The process of constructing objective (1) or (2) leads to **model constructions**, which refine graph layouts, alter topologies of models, or enhance domain-specific tasks based on semantics obtained from embeddings [PKL*18, PCZ*21, RPSM22]. If we combine and investigate objectives (1) and (3) while considering objective (2), this leads to **model debugging/quality/bias control**, e.g., graph bias or text bias examples. Last but not least, for various domains with multimodal heterogeneous datasets, **model results representation** makes all the tasks above possible.

5.7. Visualization Aspects

Classifying the visual aspects of the visualizations from the surveyed publications provides the same challenge encountered in Section 5.6, since we want to specifically target the visual aspects regarding the embeddings, and not necessarily regarding the full applications. In addition to this, in most cases, it is not straightforward to put a clear demarcation line for which data is directly related to the embeddings, and which data is not. Loosely speaking, we have chosen to include all visual components that directly aim to show information from the embedding vectors (e.g., feature values, feature value distributions, and dimensionality reduction results) or data derived from direct vector calculations (e.g., distance calculations or similarity calculations). Furthermore, since visualizing embedding data is far from a standardized task (and the level of variation is high within our surveyed corpus), we made a somewhat controversial decision to avoid detailed categories corresponding to particular metaphors or techniques here. Instead, we only focus on the categorical, binary question, and cover further details within the free-text notes. All in all, this opens for some level of ambiguity within the classification, but it nevertheless provides enough rigor for our purposes and allows us to draw some interesting general conclusions.

Our first main observation is that roughly 83% of the applications (102 entries) contain **explicit embedding representations**. This is not surprising in itself (with regard to the profile of the survey), but it is still a clear indication that using embeddings as "black boxes" is not a suitable strategy for many cases. Our second main observation is that the most dominant visual representation (used in 64 visualizations) is a scatter plot showing the 2D projection of the embeddings by using a dimensionality reduction algorithm, such as t-distributed stochastic neighbor embedding (t-SNE) or uniform manifold approximation and projection (UMAP) (Figure 5).

The main scenario for using such 2D projections is to provide the user with information on how the data points are distributed and/or how they are grouped/clustered in the embedding space with regard to pairwise similarity [HPvU*18,ZSHL18,HPX*21,BCH*22]. The conformity of the design of these 2D projections and the abundance of their use lead us to conclude that this is a *de-facto* standard for visualizing embedding vectors. As mentioned, the variation is much higher for other visual representations, with far fewer occurrences for each specific type. Hence, a main takeaway from this section is the observation that (apart from the 2D projections) the level of standardization of visualizations for embeddings is low, and that many applications contain some element(s) of unique design (Figure 6). The following are examples of some reoccurring visualizations: word clouds (8 entries) [CKC19,CBS*19,WMW*22], node-link diagrams (13 entries) [FZCM20,HKD*21,CGH*22] (see Figure 7), and scatter plots with other content than 2D projections [LCSEK19]. To directly view all the existing techniques using a specific visual representation, e.g., word clouds, readers can try typing "word cloud" in the search bar of our visual browser (cf. Figure 4)—assuming that the respective term was mentioned in free-text annotation notes.

Among many other techniques that create a typical scatter plot and assign the color of 2D projections based on either the clusters each point belongs to or the labels of those points, we identify several other options to assign coloring, such as average distance in the embedding space [BZQ*21]. Moreover, in VA systems such as *Emblaze* [SWP22] and the tool from Heimerl et al. [HG18], users are able to customize color assignments to be able to compare different embedding spaces or data co-occurrences in the same view. There are also a few techniques augmenting the scatter plot by constructing a pattern graph and Voronoi maps for each cluster subspace via Delaunay triangulation [LDL*20, VMZL22]. *AnchorViz* chooses to use non-orthogonal layouts to project the embedding space [SGR*20], and the VA technique from Heimerl et al. allows users to define the axis of the 2D projection [HG18].

5.8. Interaction Aspects

Going from the visual aspects to the interaction techniques, we once again want to underline that we specifically target the embeddings, and not the full applications. This is the explanation as to why we have only 77 entries for the category **interaction techniques support** [FCH*22, GZRP*22, RMH*22] (as compared to the 102 entries with **explicit embedding representation** in an interface), which in turn gives that roughly 25% of the applications that specifically show some embedding data do not allow any direct interaction with it [CKC19,BZQ*21,WMW*22]. Since interacting with embedding data is a non-standardized task (cf. Section 5.7), we decided to also approach this aspect with a single nominal category + free-text details, rather than introducing a variety of individual interaction categories. Browsing through the classification result, we can find many of the common interaction techniques, such as *click & select*, *pan & zoom*, *search & filter*, *details-on-demand*, etc., as well as customized variants.

In our free-text details of annotations, we found several interesting interaction techniques unique to embeddings, besides other more common user interactions, such as selecting a point on a scatter plot to show details on demand. Section 5.6 discusses three primary analytical tasks for embedding spaces. Viewing local neighborhoods is one of them. While common techniques are using color assignment or hover to show subsets of embedded items, techniques like *Cytosplore* [HPvU*16] allow users to refine and merge multiple clusters by clicking on one or more clusters represented by a heatmap with color representing the homogeneity. Another way to introduce a more customized display of embedding-derived views is by enabling user-defined filtering. In *Urban Rhapsody* [RMH*22], users can stack projections and re-project a subset of data they choose to keep. User-defined inclusion thresholds for data and cluster filtering vary greatly by each VA technique.

While some approaches enable users to filter out data points based on their direct observations and perceptions, others incorporate additional calculations to provide users with separate metric views to reference, or impose constraints to limit the range of items to be filtered [PDD*22, GZRP*22]. For example, in *CriPAV* [GZRP*22], a probability-based linear selection mechanism is implemented for filtering a set of embedded items as an alternative to brushing. Here, they use parameters to define the slope and positions of a straight line that divides points in the scatter plot into two different groups based on a function that calculates the probability and/or intensity of embedded items. Therefore, the visual representation will adapt accordingly to changes in the underlying data, as well as any user-defined changes in aggregation scores or metrics (Figure 6(a)) [SKB*22].

5.9. Evaluation Aspects

The type of evaluation that was performed is the last part of the classification scheme used for this survey. We separate between **evaluation of computational components** (53 cases) [HMW*15, ZDL21, AAM*21] and **evaluation of visual/interactive components** (84 cases) [MDL07, BMS17, SH20, PdSP*22], and also allow for free-text specification of the evaluation details. Going through the results, we find the same examples of evaluation methods (e.g., expert reviews, user studies, or questionnaires) as would be expected for evaluating any general visualization systems.

The free-text specifications for all papers in our corpus reveal that the attributes most commonly emphasized by authors are "usefulness" (35 entries), followed by "effectiveness" (25 entries), "accuracy" (22), "performance" (14), "usability" (15), "quality" (11), "scalability" (7), "efficiency" (6), "stability" (2), and "speed" (2), along with other terms such as "helpfulness", "capability", "validity", "consistency", or "sensitivity", each occurring only once. For each entry, those free-text details can be found under "evaluation details" in the survey browser, and we conducted a basic statistical

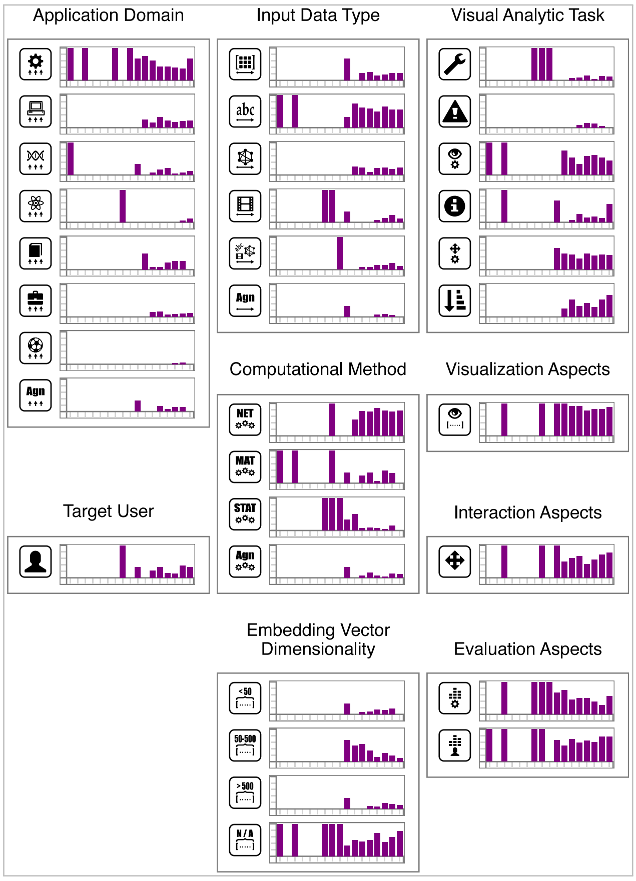

Figure 8: *Sparklines representing the support for each category within our survey data set (see Table 1 for the legend) relative to the total number of entries for the same year. (Larger version included in supplementary materials)*

frequency count of those annotations to generate the results above. Based on the results and observations during our annotation process, it is evident that only a few papers write down their considerations for the stability and robustness of the embedding methods, and not all the papers evaluate the computational component in the VA pipeline compared to user studies performed for visual components. Additional discussion related to the stability of embedding methods is provided in Section 7.1.4.

6. Survey Data Analyses

In this section, we continue analyzing the collected survey data, albeit with the overall trends and patterns in mind rather than focusing on individual example approaches.

Temporal Trends While we have presented the overall support for individual categories in our data in Table 1 and overall temporal distributions of the survey entries in Figure 3, particular categories might have become more or less prominent over time. We can thus consider the temporal distribution for each category (normalized per respective year), as displayed in Figure 8. From these

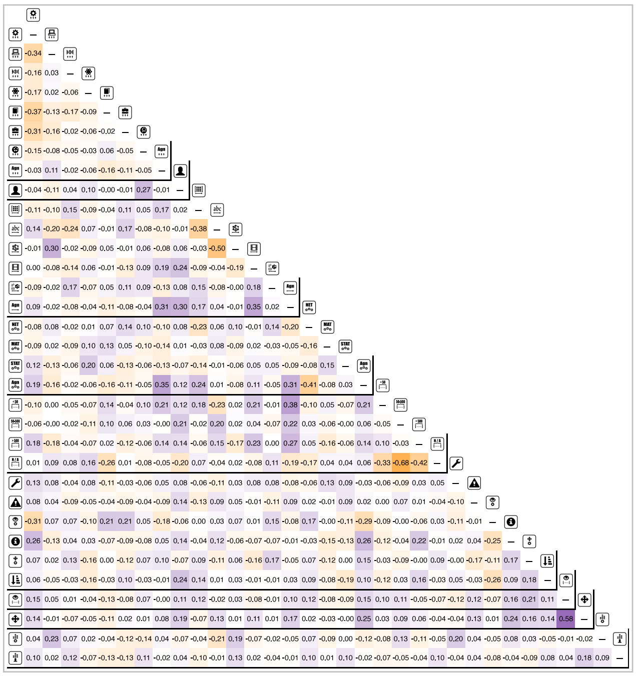

Figure 9: *Correlation matrix for the categories based on the survey data set. (Larger version included in supplementary materials)*

results, we can notice the stable interest for domain applications in ML/AI and text data, for instance, which could be expected based on the strong support for these categories in the survey data. A more interesting result here can be observed with respect to the computational methods, where the entries included in the past several years demonstrate strong interest for neural network approaches for embedding calculations. This result makes sense, given the existing trends in NLP/ML/AI, but it also confirms the same trend for the more particular scope of our survey (including the visual analytic perspective).

Category Correlations Another perspective to consider for the categorized data is whether the categories tend to co-occur in particular patterns, or even whether the use of some particular category typically means that some other category would not be supported by the respective visual analytic approach. To address this question, we have conducted correlation analysis of the survey data and computed Pearson's r coefficient values for pairs of categories. The results are presented in Figure 9, with the shades of indigo blue indicating positive correlation, while orange indicates negative correlation values. Focusing on the top positive values, some of the interesting findings here include correlation between the application domain of computing and graph input data ($r = 0.30$); support for embedding vectors under 50 dimensions and data-agnostic approaches ($r = 0.38$); and description of the target user for input data type-agnostic approaches ($r = 0.30$). We could also expect to find the rather strong (but interestingly enough, far from absolute) correlation between the visual representation and interaction categories ($r = 0.58$, the largest positive correlation value currently). Positive correla-

Table 2: *Topic modeling results for the underlying publication abstracts (computed with BERTopic [Gro22]).*

Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Outliers
DATA VISUAL USERS EMBEDDINGS VISUALIZATION EMBEDDING MODELS INTERACTIVE USER MODEL	GRAPH EMBEDDING NETWORKS NODE NODES NETWORK GRAPHS STRUCTURAL COMPARISON VISUAL	REVIEWS TOPIC ANALYTICS DATA SOCIAL SENTIMENT EVENTS VISUAL ANALYSIS GROUPS	URBAN MOBILITY PATTERNS CRIME TRANSIT AREAS NOISE CITY APPROACH DATA	SEQUENTIAL SEQUENCES DETECTION EXECUTIONS ANOMALY ANOMALOUS MOOC TIME RARE BEHAVIORS	ROLES INFORMAL ANALYSIS MATCHES CHARTS ANALYSTS TACTICS TENNIS CONSTELLATIONS TABLE	DATA ML VISUAL GESTURE TEAMS PAINTING ANALYSIS DEPENDENCIES KERNEL USER
65 entries, including: [LBT*18] [KCK*19] [WHC*23] [MWJ22] [WJM*22]	14 entries, including: [LNH*18] [CGH*22] [CZC*17] [XTYL20] [LHZ*22]	14 entries, including: [HKD*21] [XTYL18] [CKN22] [AYL*22] [CKC20]	8 entries, including: [ZDL21] [BZQ*21] [MHL*20] [RMH*22] [GZRP*22]	5 entries, including: [XXM19] [LDL*20] [DSS*23] [MDL07] [WLHO19]	4 entries, including: [WWC*21] [KW18] [XWX*22] [XBL*18]	12 entries, including: [FCH*22] [LTHL21] [SSKEA21] [ZWW*22] [BPP*16]

tions between domain-agnostic vs computation-agnostic ($r = 0.35$), domain-agnostic vs data-agnostic ($r = 0.31$), and computation-agnostic vs data-agnostic ($r = 0.31$) categories are also reasonable to expect. Other noteworthy results are the correlation between the model results explanation task and the AI/ML domain ($r = 0.26$), and the same task with computation-agnostic approaches ($r = 0.26$). With respect to negative correlations, we mainly find "competition" within the same aspects, such as lack of explicit embedding dimensionality description and the support for embedding vectors between 50–500 dimensions ($r = -0.68$, the strongest negative result currently); however, the negative result for text and graph data ($r = -0.50$) is interesting, indicating the more focused application of visual analytic approaches, perhaps. Regarding the visual analytic tasks, model results representation is negatively correlated with model results explanation ($r = -0.25$) and comparison/selection ($r = -0.26$). Additionally, the model results representation task is, interestingly enough, negatively correlated with the AI/ML domain ($r = -0.34$) in the current survey data, which is related to the support/correlation for other VA tasks with that domain.

Frequent Category Co-occurence Patterns While Pearson's correlation analysis allows us to investigate linear relationships between categories in the survey data, and other correlation analysis methods could reveal non-linear relationships, they are limited to *pairs* of categories. In order to investigate patterns involving more than two categories, we have applied frequent pattern mining using the FPGrowth [HPY00] algorithm. According to the respective results and focusing on largest pattern sizes (i.e., sets of co-occurring categories), we can describe the following profile of a VA technique involving embeddings: designed for the AI/ML domain, using neural network approaches for embedding computations (while not specifying explicitly the size of respective embedding vectors), supporting the visual analytic task of comparison/selection, explicit visual representation of embeddings (or derived results) and interaction with them, while discussing evaluation of such human-centered aspects in the respective publication. This pattern is supported by 9 entries (7% of the complete data set) in our current survey data [FZCM20, GHM21, SKB*21, CZG*22, PdSP*22, RSL*22, ZJQH22, HHS*23, WHC*23]. Shorter patterns occur more frequently, up to 15 entries (12%) supporting the pattern of the AI/ML domain, text data, neural networks, and interactive visualization & evaluation concerns . The complete list of the category patterns of length 7 and 6 (with respective citation keys) is included in supplementary materials.

Table 3: *Authorship count distribution. The current data set includes **122** entries corresponding to visual analytic approaches/studies and **503** authors in total.*

#entries	1	2	3	4	5
#authors	428	59	10	3	3

Topics in Publication Abstracts Switching from the annotated category data, we have also conducted topic modeling with the abstracts for the base publications from our survey. The motivation for this analysis is to provide an alternative way of establishing groups of publications/approaches with respect to the descriptions provided by the original publication authors, rather than the survey data annotation conducted as part of this STAR—and, furthermore, to identify salient key terms that describe such groups of publications. Specifically, we have conducted topic analysis using BERTopic [Gro22] for the current survey data: this approach essentially assigns documents into one of clusters (or labels the document as an outlier) and afterwards identifies descriptive terms for each cluster. The results from this analysis are presented in Table 2, including the top 10 terms and up to 5 most strongly associated entry references for each topic. It is noteworthy that the largest topic/cluster includes 65 publications (53% of our survey data set), but the set of respective terms such as "data", "visual", "users", and "embeddings" provides a reasonable interpretation of this broad group. Other topics include terms more specific to graph/network data, text analytics, urban data, sequences & anomalies, and sports data; the topics thus paint a picture of the main themes addressed by the existing work on the use of embeddings in visual analytics.

Authorship Statistics Besides the contents of publications themselves, it was also interesting to consider the state of the research community contributing to this topic. While our current data set is limited to 122 publications, the current authorship statistics presented in Table 3 indicate the presence of several prominent researchers with 3–5 relevant publications. The top authors according to our current data are M. El-Assady (5 entries), K.-L. Ma (5), K. Mueller (5), N. Elmqvist (4), H. Lin (4), and Y. Tao (4).

Co-Authorship Network Furthermore, we have extracted the co-authorship network, weighted by the number of entries for author nodes and co-authorship edges and conducted further network analyses in Gephi [BHJ09] (see Figure 10). The current network includes 65 weakly connected components, and the figure highlights the top 8 components (based on the size). There is one giant com-

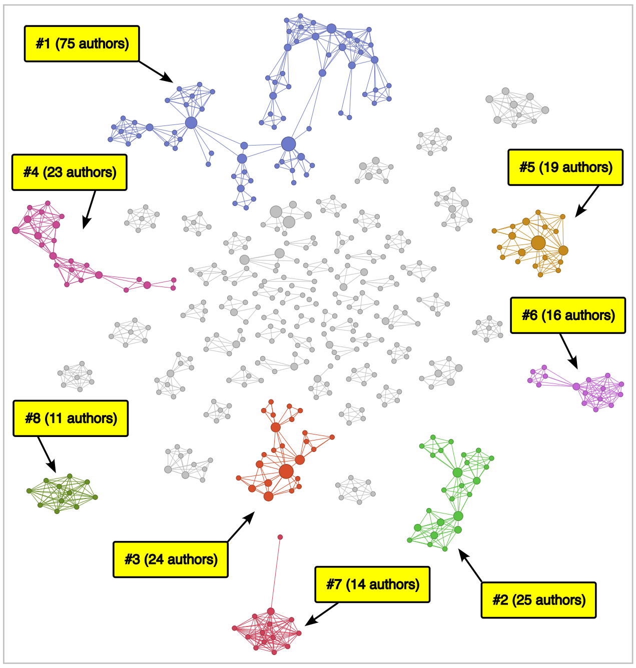

Figure 10: *Co-authorship network for the entries currently included in our survey. (Larger version with node labels included in supplementary materials)*

ponent with 75 authors that includes K. Mueller (component #1 in the figure), while components with 25 (incl. L.G. Nonato, #2), 24 (incl. M. El-Assady, #3), and 23 (incl. D.H. Chau, #4) authors follow. These results might be interesting for the readers that intend to identify and follow further work of the respective authors and research groups on this topic in the future.

7. Discussion

In this section, we bring together the insights gathered from our review of 122 papers included in this survey. We synthesize some of the findings made in the preceding sections into meta-level discoveries. Specifically, motivations, trends, concerns, and open challenges will be discussed based on two angles: embeddings for visual analytics (Section 7.1) and visual analytics for embeddings (Section 7.2). The former refers to the need to use embedding techniques in enhancing a VA system, while the latter refers to those visualization interfaces dedicated to understanding and exploring embeddings.

7.1. Embeddings for Visual Analytics (embeddings4VA)

There are multiple motivations to incorporate embeddings as a module or as a UI panel for a visual analytics tool, and this choice may be content-dependent. As discussed in Section 5.1, the necessity to include embedding techniques arises from context-specific research questions and collaboration with domain experts. On the other hand, there are more general design considerations for incorporating a dedicated embedding view into a VA system.

The unique and specific design decisions (both context-dependent and independent) for integrating embeddings in visual analytics are (1) to create an abstraction of raw data in a vector space for information retrieval across the same or different data types/modalities; (2) to offer a semantically proximal space for expert users to perceive global data distributions; (3) to act as interaction launch pads for details-on-demand with the help of clustering algorithms; and (4) to guide and recommend exploration targets using distance metrics. Section 7.1.1 elaborates on design decisions (2) and (3) in terms of user interaction, while Section 7.1.2 focuses on the visual representations of these decisions. In Section 7.1.3, we examine the evaluations of design decision (2), and in Section 7.1.4, we discuss all four design decisions from a societal perspective.

7.1.1. Embeddings as a Basis for Navigation

Unless the main objective of a tool is to address some aspect of context-independent embedding, the embeddings are **explicitly visually represented** in a UI panel, because it helps to "evaluate" the primary visual analytical tasks [SCR*23]. "Evaluation" here refers to users using such a panel to perceive the performance of a VA system, specifically, whether the system accurately addresses the intended domain challenges. Despite other views displaying various metrics that fulfill the system's primary objective, the embedding view sheds light on the original data distributions. For instance, Li et al. propose a VA tool for analyzing medical records, which includes a distribution view to embed and project all patients' medical records into a scatter plot. It provides a general understanding of the data distribution [LYY*20]. Embeddings here are not only a numerical representation of data objects, but their visual representations also act as some sort of "ground truth" representations of the original dataset. When users perform filtering or selection on a 2D projection of the embedding space, which is assumed to be a representation that best preserves the semantic closeness of data points, adding an embedding view can make the operation more convincing by providing a direct comparison before and after the operation in a global view. This aims to enhance the credibility and legibility of the tool.

In visual analytics interfaces that aim to optimize computational models with human-in-the-loop processes, the **interaction** sequence typically begins with the embedding panel. It acts as a launch pad for a deeper exploration of subspaces with the help of clustering techniques [MM23]. Here, users can evaluate the data using their domain knowledge and perform interactions like assigning labels to neighboring points or merging objects into clusters. These interactions result in updates to other panels, which help users achieve high-level analytical goals such as topic modeling and text alignment [EAKC*20, MWJ22]. These interactions result in an update to the entire interface, with the embedding view showing global distributions that provide feedback on the user interaction, thereby enhancing the human-in-the-loop process [SJB*17, BNL*18]. For the process described above, human knowledge is incorporated, validated, and propagated from the embedding view to other linked panels.

If we consider this process as active involvement in a VA system, then a more passive involvement would be when the embeddings directly provide insights to users. Computing embeddings is part of the process for recommending similar objects to a target user input, such as keyword suggestions or query search as mentioned in Section 5.1.2. In such cases, users receive knowledge from the embedding view, obtain serendipitous findings, and incorporate them to achieve better performance in the analytical goals.

7.1.2. Embeddings for Visual Channel Assignments

Visualizing embeddings explicitly can be a way to augment the current panel with additional information and features. It is a common technique to position visual elements based on the results from dimensionality reduction or multi-dimensional scaling methods of the corresponding embeddings. This process is referred to as *spatialization* by various authors [ERT*17]. Regardless of the types of data being analyzed, proximity in the embedding space conveys information. The closeness, separability, and density between visual objects may indicate the presence of inherent similarities or other types of relationships between them. However, it remains an open question whether users can perceive a coherent message from such 2D embedding projections.

Visual variable effectiveness is linked to the level of perceptual precision, which refers to the accuracy of values being interpreted when encoded by visual channels [BCF20]. As mentioned in Section 5.9, most papers in our corpus evaluate effectiveness, usefulness, and other system performances by conducting user studies. However, evaluating embedding-related views as a whole does not necessarily mean all of the visual encodings work in an effective way. Using embeddings to assign the position of visual elements is often data- and model-dependent.

Even though the generated high-dimensional embedding vectors always manifest some latent patterns, a 2D summarization and position assignment may not always result in a clear visual separation between groups of data. The designers' intent for assigning embedding-derived positions to data objects is often that the users should be able to locate clusters. Such intention may not always lead to accurate perceptions, especially when there are fewer observable between-cluster separations. Users may falsely find patterns that are actually statistically insignificant. Thus, it is questionable whether the users' interpretation of such visual representations is reliable and functional.

To address this issue, many works add additional guidance to improve separation. For instance, in *Semantic Concept Spaces* [EAKC*20], Voronoi tessellation is used to enhance the visual association of words by adding concept boundaries in the embedding space. Even though there exist multiple other types of layout enrichment, as discussed in Nonato et al.'s survey about multidimensional projection techniques [NA19], in our corpus, we observe much less variety in applying enrichment.

7.1.3. Learnability, Stroop Effect, and "Slow Analytics"

Depending on the **target user**, some papers in our corpus **evaluate** whether their proposed system is easy to learn [XBL*18]. Practicality [FCH*22], learnability [JSR*19], and readability [PCZ*21] are explicitly accessed for a small set of papers, while most other papers' primary concern is showing the usefulness of their tool.

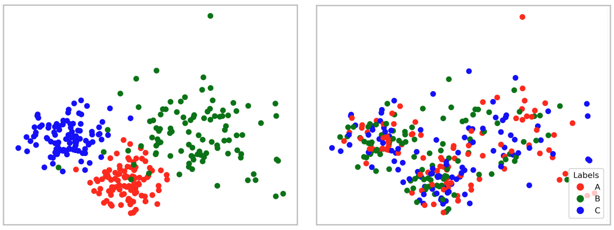

(a) (b)

Figure 11: *An example of two types of color encoding for a 2D scatter plot: (a) applying double encoding with colors to represent identified clusters, and (b) applying color encoding based on another labeling/attribute. The latter case demonstrates the potential issues with visual encoding of multivariate data when making use of DR and projection plots.*

Evaluating a VA system's learnability, effectiveness, and usefulness are deeply intertwined. Nevertheless, the multivariate nature of embedding-related projection panels produces unique challenges in handling learning difficulties. By multivariate, we mean a panel that projects data embeddings into a 2D space while the data incorporate additional attributes and labeling. Visually, this means that users can observe clusters of data points based on their similarities in a high-dimensional embedding space, but the points in the same cluster may also share different labeling.

Figure 11(a) illustrates the effectiveness of double encoding in distinguishing between clusters by conveying the same information through both position and color, but as Figure 11(b) shows, assigning colors based on a different grouping can alter the perception. For instance, in *ConceptVector* [PKL*18], the word cluster view uses position and coloring as two main visual channels, where word embeddings determine the position, and the coloring is another grouping label applied to the data. Such design may create incongruence between the encoded positions and colors, similar to the Stroop effect [Str35] (individuals may find reading a word printed in blue font as "red" difficult). Moreover, this issue makes adding additional annotations to the embedding view difficult.

Encoding additional information can result in a fruitful process for addressing analytical tasks, as Lupi's Data Humanism manifesto and "slow analytics" movement from Bradley et al. encourage viewers to spend time with the data to ensure those data are retained [BEF17]. On the other side of the spectrum, many other works choose dual-coding [WWC*21, LZ23] as shown in Figure 11. There are also many approaches that simply list items using embedding-derived similarity measures in a table, with no additional visual representation. The contrast of designing decisions in our corpus reveals the role embedding as a visual representation plays within a VA system. With multiple visual channel assignments, the former often plays a critical role in the human-in-the-loop processes. Meanwhile, the latter that uses dual encoding often tries to use embedding as a supplementary panel. It provides a reference to the original data distribution to add trustworthiness to the VA interface.

7.1.4. Stability, Reproducibility, and Bias

If we review a typical VA pipeline that produces **explicit visual representations** of embeddings, we would likely find that the embeddings are never capable of representing the data fully accurately. Nevertheless, many papers in our corpus may have a slightly different assumption, as stated in Section 7.1.1. In the previous Section 7.1.2, we discussed how different data distributions may produce either easy-to-separate clusters or borderline cases that may be misleading in falsely identifying statistically insignificant groupings. Besides the varying distribution of source data that may influence the reliability of a VA system, two additional factors further hinder the credibility of an interface: the absence of stability and the presence of societal bias.

The NLP community has produced a multitude of papers examining the issue of stability when various computational models generate embeddings. Embedding-based word similarities, especially nearest-neighbor distances, are highly sensitive to small changes in the training corpus [AM18], yet some VA tools in our corpus rely on embedding-generated keywords as feedback guiding users' interactions.

Multiple other factors, from data properties to algorithm properties, influence the stability of embeddings [WKM18]. Hyperparameter choices may also influence the generation process [BBA21]. However, according to Section 5.5, not all papers give an exact number in the output embedding dimensions. Even a smaller set of papers mention how their embedding model is trained, fine-tuned, and evaluated. The **interactivity** in visualization offers a unique way to address those concerns of instability, which is to let users decide. We observe papers state that their visualization is model-independent, offering the possibility to switch between visualizations and adding input features that allow users to decide on certain hyperparameters [SH20, SKB*22, SSJ*22]. Still, for many VA tools that aim at addressing a domain-specific problem and providing insights to its viewers, the need for more specification for computational components, together with the hidden instability of the chosen computational method, would hinder the reproducibility of their work.

Visualization tools seek to be objective in delivering information. Societal bias may seem to be less of a concern for tools we labeled as **domain-agnostic** in Section 5.1.8, where they intend to offer visual exploration of user-defined embeddings. However, under a different context, the bias encoded in embeddings would be detrimental to the credibility of a VA system. Especially for those downstream applications we categorize under the domain of **humanities, social sciences, and education**, and **business, management, governance, law**, they intend to assist in the decision-making process for governments, companies, and domain specialists. As Baumer et al. argue in their recent study [BJSM22], further perspectives beyond the purely analytical one apply in such scenarios (namely, the political one), which must be taken into account when considering the existing or designing new such approaches.

Understanding how social biases are reflected in word embeddings have been widely studied in NLP and CL communities, and bias can be encoded in multiple social dimensions [JM20]. Even though many visual analytical tools use word embeddings to represent data in a VA pipeline for decision-making, they fail to include reflections and considerations to mitigate bias. Incorporating **embedding-related visual representations** sometimes aims to empower the VA system with abilities to navigate through high-dimensional space. At the same time, such visual representations amplify the existence of biases by making recommendations to the users and visually displaying connections between words that potentially would reinforce social stereotypes.

7.1.5. Contradictions, Challenges, and Opportunities

In summary, we observe two contradictions when using embedding techniques for visual analytic tools. First, designers intend to incorporate embedding-related panels to add credibility to their tool, as the **embedding view** often provides a reference to the underlying data distributions. However, introducing embeddings also brings more uncertainty. These techniques are prone to instability and often encode societal biases. Meanwhile, authors of related papers often provide limited reasoning and detail regarding the choices and decisions they make. Second, designers seek to add multiple visual channels to encode additional information in an embedding projection view. However, adding more visual encodings would influence the systems' learnability, but at the same time, one could argue that spending more time in the panel would help data be retained and ultimately perform better in high-level analytical tasks. In addition, complexity in visual encoding schemes offers more space for user interactions for many human-in-the-loop model refinement processes. There is no perfect solution to address those two contradictions, and there are always trade-offs and balances between two sides of the spectrum. Using embedding techniques to facilitate high-level tasks in visual analytic tools (embeddings for visualization) presents many open challenges and opportunities to develop community-wise suggestions and guidelines.

7.2. Visual Analytics for Embeddings (VA4embeddings)

As discussed in Section 5.1.8, **visualization** can help reveal social bias inside embedding spaces [GHM21], interpret user-defined embeddings, and compare multiple embedding spaces. Those spaces are generally high-dimensional. Visualizing them gives users a direct perception of the global geometry and potential substructures such as clusters [BCS22]. Along with many other tasks that reveal the characteristics of embeddings, visualization can communicate insights beyond gathering metrics. The interactivity within a visual analytics system allows users to build the mental map between embedding space and its underlying data distributions. In addition, it is easier to make visual comparisons across different embedding spaces.

7.2.1. Comparisons and Aggregations

Typically, **comparisons** are made between two embedding spaces in a juxtaposition fashion. Those side-by-side 2D-projection panels assist users in observing the topological differences [BCS22]. For various other works in our corpus, they provide analysis across multiple embedding spaces, as well as offer visual detection towards temporal semantic changes. Comparisons are facilitated by introducing new visual representations [XTYL20, SKB*22, XWX*22, SWP22], graph drawing techniques

[JCS*21], and computational methods [CZC*17]. Many works express changes in comparison by adding traces and lines in a scatter plot connecting different data objects or areas. Distance metrics and neighborhood comparisons are often encoded into visual representations to quantify differences. Even though many of those tools provide unique and effective ways to compare multiple embedding spaces, finding out what to compare remains challenging. For semantic changes in word embeddings [SLN*19], diachronic evolution of node embeddings [XTYL20], or any other collections of embedding spaces, the reference point for comparison and the size of aggregated embedding collections (i.e., length of snapshots) matter. As a result, one open research question would be providing linked views with other aspects of data that give users insights into selecting configurations such as reference points.

Other important questions arise when we compare a collection of embedding spaces: how would we aggregate a set of embeddings, and how might we perform comparisons analysis over a combination of embedding technology and ensemble methods? For the former question, Cavallari et al. [CZC*17] propose an inspiring way that treats embeddings not as vectors but as distributions in the low-dimensional space. Thus, one may visually analyze the embeddings of communities rather than the embeddings of individual nodes. According to our corpus, there is a limited number of papers mentioning this direction. This work leaves out open challenges in visually comparing aggregated embedding spaces. For the latter question, Witschard et al. [WJM*22] provide interfaces for comparing the performance in a combination of various embedding techniques in an ensemble learning setting. Their work intends to improve the quality of certain similarity calculations. However, one could apply a variation of such an idea to visually analyze combinations of sentence embeddings, locally-aggregated embeddings, etc.

7.2.2. Multimodality and Multi-level Analysis

In our corpus, the objects to be embedded are diverse in their data type. As mentioned in Section 5.3, besides embedding textual, numerical, and graph data, some VA tools use different neural encoders to encode data such as time-series [LDL*20, LPH*20, GZRP*22], audio [RMH*22], or images [GGW22, PDD*22]. The embedding view of some included papers may incorporate **more than one** type of data. For instance, Cabrera et al. [CDHP21] combine images and texts into the 2D projection view. For some other papers, computing metrics or embedding-related views may require more than one type of data as well. For example, many papers under the **Urban Visual Analytics** subcategory (as mentioned in Section 5.1.4) require spatial-temporal data as input for further analysis. However, many of them choose to only embed one type of data. The resulting embeddings are visually encoded and superimposed into a panel.

From the computational side of view, there has been a growing trend toward developing techniques for generating joint embeddings from heterogeneous sources of multimodal data [ZYHD20]. However, for the papers in our corpus, even though many downstream domain applications naturally obtain data from various sources, only a small number of papers in our corpus try to embed different types of data into a shared latent space. Additionally, even though multimodal techniques are applied, such as *LeSSS* [HMW*15], the resulting **visual representation** is often limited to a 2D projection and need more interactivity.

As a result, multiple gaps exist between the domain-specific diverse data sources, the advances in computational techniques, and the rather limited ways for a VA system to sufficiently represent all the information. Besides incorporating state-of-the-art embedding techniques, when designing a VA interface, there is much room for further investigation on how to encode multiple data types, build connections between them, and provide context and details when users navigate through such embedding spaces.

There exists a lot of diversity in input data. At the same time, there are multiple levels of understanding regarding the embedding space. Given word embeddings, one could analyze the representative keywords, descriptors, topics, concepts, and document-level characteristics [EAKC*20]. There is a hierarchy of analysis for the embedding space, as well as a hierarchy for user explorations. Given a large-scale data set, computing the embeddings of everything may be costly. If a higher-level embedding already shows all features of interest, it is unnecessary to generate a complete representation of the entire data set [HPvU*18]. For both cases, users need a sufficient amount of guidance. When exploring an embedding space, it is necessary to provide an overview of the current state of exploration. In addition, users need to be directed to unexplored areas based on the potential insights from analyzing its lower-level embeddings.

In general, for those VA tools aimed at understanding and exploring a large-scale embedding space (visual analytics for embeddings), it can be challenging to provide adequate guidance to users, which encompasses not only suggesting levels of analysis but also highlighting potential areas of interest.

7.3. Limitations of the Conducted Survey

One limitation of this study was the number of paper entries included for the current corpus and the small number of papers for each category, especially for application domains. We screened through approximately 1,704 papers but only included 122 papers that matched the criteria for inclusion in the analysis. As we observed, the top few reasons for those papers to be not matching our criteria include the following:

- Papers that have little interactivity or limited contributions in visualization. Since our search is not restricted to visualization-specific journals, we searched through those papers working on the computational side of embeddings, which contains the word "visualization" and its related concepts. Often, they provide a static visualization plot for demonstrating and comparing the results of their proposed embedding process.
- Papers in which the resulting embedding space is 2-dimensional or unclear. Common dimensionality reduction techniques, such as locally linear embedding (LLE) or t-distributed stochastic neighbor embedding (t-SNE), include the term "embedding" in their titles. In addition, many authors would refer to "embedding" as purely the process of using DR to map their data into 2D projections. There exist some other papers that only mention multidimensional scaling projection (MDS) in their methods, but with little computational details.

- VIS4ML papers that analyze model layers and attention mechanisms, but do not contain specific analysis of the embeddings.
- Visualization papers that use the word "embedding" to express they "embed" a visual representation or a panel to the interface.
- Visualization papers mentioning "embedding" when they explain how their ML pipeline would work.
- Papers mentioning "embedding" as one of the future works.

The semantics of the term "embedding" exceeds the definition used in this survey. Nevertheless, much of its use in the VA context is closely aligned with how other computational domains, such as NLP, would describe it. In addition, many visualization papers use word2vec, GloVe, BERT, etc., to create "embeddings" during their VA pipeline, but do not explicitly mention the term "embedding". Therefore, it is possible that much more diversity is not being included. It is expected that there is a need to query more papers and perform analysis at a larger scale with an extensive amount of keywords.

7.4. Open Challenges and Future Perspectives

This survey covers a unique and special set of papers. Analyzing them from different perspectives could result in different findings. Domain-specific and context-dependent challenges are discussed in Section 5.1. In Sections 7.1 and 7.2, we group papers by different criteria (Embeddings for VA, and VA for Embeddings), and discuss high-level visual analytical challenges from those aspects. Here, to synthesize all the discussions, we present open challenges and future perspectives based on our proposed categorization as shown in Table 1.

Input Data Type Current VA+embeddings works are primarily focused on standard word embeddings, node embeddings, etc. At the same time, most of those VA tools aim to solve complex real-world scenarios that produce **a variety of data types** from heterogeneous sources. As argued in Section 7.2, future work may target an ensemble approach of incorporating different types of embeddings (e.g., sentence embeddings rather than word embeddings) and learning joint embeddings with multimodal data. Another challenge is determining the level of generalizability a VA tool should aim for. Constraints on the input data format can be one key factor influencing generalizability, as papers in our corpus either allow user-defined inputs or build the entire system based on a fixed dataset.

Computational Method It is never an easy process to decide what to embed. Different domain-specific research questions lead to various representation challenges. Finding suitable ways to transform task descriptions into feature vectors would impact the performance of using a specific embedding technique for context-dependent tasks. Furthermore, some VA tools seek to construct a real-time interface or generate representations of large-scale data sets. As a result, one needs to identify potential bottlenecks in the computational pipeline and optimize the process for embedding generation, similarity search, and metrics calculation. Additionally, the state-of-the-art embedding methods are changing rapidly, but not a lot of VA tools in our corpus consider incorporating the most up-to-date embedding techniques and applying novel metrics beyond measuring the Euclidean distances between data objects in an embedding space. Regardless of what computational methods are employed, designers of those VA tools need to provide more justifications for their choices of embedding models and embedding-related metrics to improve the **explainability** of their work.

Visualization Aspects As discussed in Section 7.1, ensuring the credibility and trustworthiness of **visual representations for embeddings** is a challenging topic. One way to address this issue is to carefully consider what elements to present and what additional metrics to incorporate into the VA interface. For instance, displaying a **comparison** between pre- and post-embedding metrics may help add transparency to the process. Creating multiple independent views that reflect features in different dimensions of the embedding space may also help provide more aspects of measuring trustworthiness.

2D projections are the predominant visual representation of embeddings. They are arguably intuitive (at least on the surface level) and useful for many situations. However, if we intend to compare multiple embedding spaces or perform complex human-in-the-loop interactions, they may not be the optimal choice. Thus, future work could explore more diverse and novel visual representations.

Interaction Aspects One important motivation for computing embeddings is to provide suggestions that guide user **interactions**. As stated in the last paragraph of Section 5.1.4, calculating distances and neighborhoods in the embedding space can help users **navigate** through decision spaces. In our corpus, most papers related to geospatial data visualization use embeddings in such a way. Potentially, embedding-based interaction recommendations can be applied to a broader set of VA systems.

Miscellaneous Aspects There are many more interesting, yet underdeveloped topics that fit within the scope of our survey. Two areas we would like to highlight are visually exploring diachronic changes in the embedding spaces and designing novel visual interfaces for multimodal embeddings. We discussed both of them in detail under Section 7.2.

Last but not least, we believe the idea of reproducibility [FF20] ought to be emphasized for both the VA system implementations and the corresponding papers. In the context of VA+embeddings, ensuring reproducibility requires providing justifications and details on several aspects, including (1) the target audiences and usage context, (2) the embedding methods and evaluation metrics, (3) the dimensionality of the output embedding vectors, (4) the projection process, (5) the distance metrics and criteria for creating the visual encoding of semantically similar objects, and (6) the stability of the embedding method and its impact on the visual layout. In addition, it will be beneficial for the whole visualization community if authors are willing to make their tools open source. In the provided survey browser, users can get a list of open-source visualization tools in our corpus available on GitHub, for instance, by typing "github" in the search bar. More importantly, adding sufficient background information (i.e., target audiences), as well as computational details for embedding generation in the main text or as supplementary materials will greatly improve the reproducibility of a VA system.

Future Work for Surveys and Empirical Studies Here are a few research topics for future work in this survey and as potential topics for other empirical studies. First, extending our categorization to include reproducibility levels for each annotated paper will benefit those seeking to create derivations of existing works by guiding them to find suitable candidates. Second, as stated in Section 7.2.2, a more focused survey on visual analytics tools with multimodal data would be of interest and value to the visualization research community and similar audiences of this survey paper.

8. Conclusions

In this survey article, we scrutinize the existing work in visual analytics (and, to some extent, other fields) that makes use of data embedding approaches. We define the scope of this study, the corresponding inclusion/exclusion criteria, and the literature search methodology that results in 122 included survey entries based on peer-reviewed publications within and beyond the visualization venues. We propose a categorization that focuses on the domain application, embedding computation, and human-centered aspects of the visual analytic pipeline. By discussing the results of survey data annotation for this categorization as well as further survey data analyses, we are able to summarize the current state of the art on the use of embedding approaches in visual analytics, including the open challenges. Furthermore, we provide access to the categorized data within an online survey browser, which can be helpful for researchers, practitioners, and students of visualization as well as other domains and disciplines, especially considering the rising interest for this topic during the several past years and the opportunities for further novel and important contributions.

Acknowledgments

This work was partially supported through (1) the ELLIIT environment for strategic research in Sweden and (2) the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. We are also thankful to the authors of original tools who either made their code or implementation available, or provided us with new original screenshots to be used as example figures for this manuscript.

References

[AAM*21] ANDREADIS S., ANTZOULATOS G., MAVROPOULOS T., GIANNAKERIS P., TZIONIS G., PANTELIDIS N., IOANNIDIS K., KARAKOSTAS A., GIALAMPOUKIDIS I., VROCHIDIS S., KOMPATSIARIS I.: A social media analytics platform visualising the spread of COVID-19 in Italy via exploitation of automatically geotagged tweets. *Online Social Networks and Media 23* (May 2021), 100134. doi:10.1016/j.osnem.2021.100134. 1, 11, 16, 18

[AES05] AMAR R. A., EAGAN J., STASKO J. T.: Low-level components of analytic activity in information visualization. In *Proceedings of the IEEE Symposium on Information Visualization* (2005), InfoVis '05, IEEE, pp. 111–117. doi:10.1109/INFVIS.2005.1532136. 17

[AGH*23] AFZAL S., GHANI S., HITTAWE M. M., RASHID S. F., KNIO O. M., HADWIGER M., HOTEIT I.: Visualization and Visual Analytics Approaches for Image and Video Datasets: A Survey. *ACM Transactions on Interactive Intelligent Systems 13*, 1 (Mar. 2023). doi:10.1145/3576935. 1

[AHH22] ANŽEL A., HEIDER D., HATTAB G.: MOVIS: A multi-omics software solution for multi-modal time-series clustering, embedding, and visualizing tasks. *Computational and Structural Biotechnology Journal 20* (2022), 1044–1055. doi:10.1016/j.csbj.2022.02.012. 10, 14

[AKZM14] AVERKIOU M., KIM V. G., ZHENG Y., MITRA N. J.: ShapeSynth: Parameterizing model collections for coupled shape exploration and synthesis. *Computer Graphics Forum 33*, 2 (May 2014), 125–134. doi:10.1111/cgf.12310. 13, 16

[AM18] ANTONIAK M., MIMNO D.: Evaluating the Stability of Embedding-based Word Similarities. *Transactions of the Association for Computational Linguistics 6* (Dec. 2018), 107–119. doi:10.1162/tacl_a_00008. 23

[AP08] ARTSTEIN R., POESIO M.: Inter-coder agreement for computational linguistics. *Computational Linguistics 34*, 4 (Dec. 2008), 555–596. doi:10.1162/coli.07-034-R2. 4

[ARCL21] ALHARBI M., ROACH M., CHEESMAN T., LARAMEE R. S.: VNLP: Visible natural language processing. *Information Visualization 20*, 4 (Oct. 2021), 245–262. doi:10.1177/14738716211038898. 8

[AS05] AMAR R. A., STASKO J. T.: Knowledge precepts for design and evaluation of information visualizations. *IEEE Transactions on Visualization and Computer Graphics 11*, 4 (July–Aug. 2005), 432–442. doi:10.1109/TVCG.2005.63. 17

[AX19] ALMEIDA F., XEXÉO G.: Word embeddings: A survey. *arXiv preprint arXiv:1901.09069* (2019). doi:10.48550/arXiv.1901.09069. 1

[AYL*22] AHN Y., YAN M., LIN Y.-R., CHUNG W.-T., HWA R.: Tribe or not? Critical inspection of group differences using TribalGram. *ACM Transactions on Interactive Intelligent Systems 12*, 1 (Mar. 2022), 1–34. doi:10.1145/3484509. 11, 20

[BBA21] BORAH A., BARMAN M. P., AWEKAR A.: Are Word Embedding Methods Stable and Should We Care About It? In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media* (2021), HT '21, ACM, pp. 45–55. doi:10.1145/3465336.3475098. 23

[BCF20] BERTINI E., CORRELL M., FRANCONERI S.: Why shouldn't all charts be scatter plots? Beyond precision-driven visualizations. In *Proceedings of the 2020 IEEE Visualization Conference* (2020), VIS '20, IEEE, pp. 206–210. doi:10.1109/VIS47514.2020.00048. 22

[BCH*22] BÄUERLE A., CABRERA Á. A., HOHMAN F., MAHER M., KOSKI D., SUAU X., BARIK T., MORITZ D.: Symphony: Composing interactive interfaces for machine learning. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2022), CHI '22, ACM, pp. 1–14. doi:10.1145/3491102.3502102. 18

[BCjC19] BASTA C., COSTA-JUSSÀ M. R., CASAS N.: Evaluating the underlying gender bias in contextualized word embeddings. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing* (Aug. 2019), GeBNLP '19, Association for Computational Linguistics, pp. 33–39. doi:10.18653/v1/W19-3805. 14

[BCS22] BOGGUST A., CARTER B., SATYANARAYAN A.: Embedding Comparator: Visualizing differences in global structure and local neighborhoods via small multiples. In *Proceedings of the 27th International Conference on Intelligent User Interfaces* (2022), IUI '22, ACM, pp. 746–766. doi:10.1145/3490099.3511122. 8, 9, 16, 23

[BCV13] BENGIO Y., COURVILLE A., VINCENT P.: Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence 35*, 8 (Aug. 2013), 1798–1828. doi:10.1109/TPAMI.2013.50. 1, 4

[BDVJ03] BENGIO Y., DUCHARME R., VINCENT P., JANVIN C.: A neural probabilistic language model. *Journal of Machine Learning Research 3* (Mar. 2003), 1137–1155. URL: https://dl.acm.org/doi/abs/10.5555/944919.944966. 3

[BEF17] BADAM S. K., ELMQVIST N., FEKETE J.: Steering the craft: UI elements and visualizations for supporting progressive visual analytics. *Computer Graphics Forum 36*, 3 (June 2017), 491–502. doi:10.1111/cgf.13205. 11, 15, 16, 22

[BHJ09] BASTIAN M., HEYMANN S., JACOMY M.: Gephi: An open source software for exploring and manipulating networks. In *Proceedings of the 3rd International AAAI Conference on Weblogs and Social Media* (2009), ICWSM '09, AAAI Press, pp. 361–362. URL: https://aaai.org/ocs/index.php/ICWSM/09/paper/view/154. 20

[BIVD22] BHATTACHARJEE K., ISLAM A., VAIDYA J., DASGUPTA A.: PRIVEE: A visual analytic workflow for proactive privacy risk inspection of open data. In *Proceedings of the 2022 IEEE Symposium on Visualization for Cyber Security* (2022), VizSec '22, IEEE, pp. 1–11. doi:10.1109/VizSec56996.2022.9941431. 9

[BJSM22] BAUMER E. P. S., JASIM M., SARVGHAD A., MAHYAR N.: Of course it's political! A critical inquiry into underemphasized dimensions in civic text visualization. *Computer Graphics Forum 41*, 3 (June 2022), 1–14. doi:10.1111/cgf.14518. 23

[BKW16] BECK F., KOCH S., WEISKOPF D.: Visual analysis and dissemination of scientific literature collections with SurVis. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (Jan. 2016), 180–189. doi:10.1109/TVCG.2015.2467757. 4

[BLB*17] BREHMER M., LEE B., BACH B., RICHE N. H., MUNZNER T.: Timelines revisited: A design space and considerations for expressive storytelling. *IEEE Transactions on Visualization and Computer Graphics 23*, 9 (Sept. 2017), 2151–2164. doi:10.1109/TVCG.2016.2614803. 11

[BM13] BREHMER M., MUNZNER T.: A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (Dec. 2013), 2376–2385. doi:10.1109/TVCG.2013.124. 17

[BMS17] BERGER M., MCDONOUGH K., SEVERSKY L. M.: cite2vec: Citation-driven document exploration via word embeddings. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (Jan. 2017), 691–700. doi:10.1109/TVCG.2016.2598667. 8, 17, 18

[BN21] BIAN Y., NORTH C.: DeepSI: Interactive deep learning for semantic interaction. In *Proceedings of the 26th International Conference on Intelligent User Interfaces* (2021), IUI '21, ACM, pp. 197–207. doi:10.1145/3397481.3450670. 8, 17

[BNL*18] BERGER M., NAGESH A., LEVINE J., SURDEANU M., ZHANG H.: Visual supervision in bootstrapped information extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (2018), EMNLP '18, ACL, pp. 2043–2053. doi:10.18653/v1/D18-1229. 8, 17, 21

[BPP*16] BARBOSA A., PAULOVICH F. V., PAIVA A., GOLDENSTEIN S., PETRONETTO F., NONATO L. G.: Visualizing and interacting with kernelized data. *IEEE Transactions on Visualization and Computer Graphics 22*, 3 (Mar. 2016), 1314–1325. doi:10.1109/TVCG.2015.2464797. 8, 17, 20

[BZQ*21] BAI J., ZHANG H., QU D., LV C., SHAO W.: FGVis: Visual analytics of human mobility patterns and urban areas based on F-GloVe. *Journal of Visualization 24*, 6 (Dec. 2021), 1319–1335. doi:10.1007/s12650-021-00775-x. 11, 15, 18, 20

[CBS*19] CHANDRASEGARAN S., BRYAN C., SHIDARA H., CHUANG T.-Y., MA K.-L.: TalkTraces: Real-time capture and visualization of verbal content in meetings. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), CHI '19, ACM, pp. 1–14. doi:10.1145/3290605.3300807. 12, 18

[CDHP21] CABRERA Á. A., DRUCK A. J., HONG J. I., PERER A.: Discovering and validating AI errors with crowdsourced failure reports. *Proceedings of the ACM on Human-Computer Interaction 5*, CSCW2 (Oct. 2021), 1–22. doi:10.1145/3479569. 8, 17, 24

[CEBV22] CHANDA A. K., EGLESTON B. L., BAI T., VUCETIC S.: MedCV: An interactive visualization system for patient cohort identification from medical claim data. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (2022), CIKM '22, ACM, pp. 4828–4832. doi:10.1145/3511808.3557157. 6, 10, 14, 17

[CGH*22] CHENG S., GIESEN J., HUANG T., LUCAS P., MUELLER K.: Identifying the skeptics and the undecided through visual cluster analysis of local network geometry. *Visual Informatics 6*, 3 (Sept. 2022), 11–22. doi:10.1016/j.visinf.2022.07.002. 11, 12, 18, 20

[CK15] CERNEA D., KERREN A.: A survey of technologies on the rise for emotion-enhanced interaction. *Journal of Visual Languages and Computing 31, Part A* (2015), 70–86. doi:10.1016/j.jvlc.2015.10.001. 4

[CKC19] CHANG Y.-C., KU C.-H., CHEN C.-H.: Social media analytics: Extracting and visualizing Hilton hotel ratings and reviews from TripAdvisor. *International Journal of Information Management 48* (Oct. 2019), 263–279. doi:10.1016/j.ijinfomgt.2017.11.001. 12, 18

[CKC20] CHANG Y.-C., KU C.-H., CHEN C.-H.: Using deep learning and visual analytics to explore hotel reviews and responses. *Tourism Management 80* (Oct. 2020), 104129. doi:10.1016/j.tourman.2020.104129. 12, 20

[CKN22] CHANG Y.-C., KU C.-H., NGUYEN D.-D. L.: Predicting aspect-based sentiment using deep learning and information visualization: The impact of COVID-19 on the airline industry. *Information & Management 59*, 2 (Mar. 2022), 103587. doi:10.1016/j.im.2021.103587. 12, 15, 20

[CMJ*20] CHATZIMPARMPAS A., MARTINS R. M., JUSUFI I., KUCHER K., ROSSI F., KERREN A.: The state of the art in enhancing trust in machine learning models with the use of visualizations. *Computer Graphics Forum 39*, 3 (June 2020), 713–756. doi:10.1111/cgf.14034. 4, 5, 6

[CMJK20] CHATZIMPARMPAS A., MARTINS R. M., JUSUFI I., KERREN A.: A survey of surveys on the use of visualization for interpreting machine learning models. *Information Visualization 19*, 3 (July 2020), 207–233. doi:10.1177/1473871620904671. 4, 7, 9

[CSH05] COZENS P. M., SAVILLE G., HILLIER D.: Crime prevention through environmental design (CPTED): A review and modern bibliography. *Property Management 23*, 5 (Jan. 2005), 328–356. doi:10.1108/02637470510631483. 11

[CTL18] CHEN J., TAO Y., LIN H.: Visual exploration and comparison of word embeddings. *Journal of Visual Languages & Computing 48* (Oct. 2018), 178–186. doi:10.1016/j.jvlc.2018.08.008. 3, 8

[CWB*11] COLLOBERT R., WESTON J., BOTTOU L., KARLEN M., KAVUKCUOGLU K., KUKSA P.: Natural language processing (almost) from scratch. *Journal of Machine Learning Research 12* (Nov. 2011), 2493–2537. 3

[CWDH09] CHEN Y., WANG L., DONG M., HUA J.: Exemplar-based visualization of large document corpus. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (Nov. 2009), 1161–1168. doi:10.1109/TVCG.2009.140. 8, 15

[CWPZ19] CUI P., WANG X., PEI J., ZHU W.: A survey on network embedding. *IEEE Transactions on Knowledge and Data Engineering 31*, 5 (May 2019), 833–852. doi:10.1109/TKDE.2018.2849727. 3, 4, 7

[CZC*17] CAVALLARI S., ZHENG V. W., CAI H., CHANG K. C.-C., CAMBRIA E.: Learning community embedding with community detection and node embedding on graphs. In *Proceedings of the 2017 ACM Conference on Information and Knowledge Management* (2017), CIKM '17, ACM, pp. 377–386. doi:10.1145/3132847.3132925. 9, 15, 16, 20, 24

[CZG*22] CHEN Y., ZHANG Q., GUAN Z., ZHAO Y., CHEN W.: GEMvis: A visual analysis method for the comparison and refinement of graph embedding models. *The Visual Computer 38*, 9 (Sept. 2022), 3449–3462. doi:10.1007/s00371-022-02548-5. 16, 20

[CZIM18] CHENG S., ZHONG W., ISAACS K. E., MUELLER K.: Visualizing the topology and data traffic of multi-dimensional torus interconnect networks. *IEEE Access 6* (Sept. 2018), 57191–57204. doi:10.1109/ACCESS.2018.2872344. 9, 10, 16

[DCLT19] DEVLIN J., CHANG M.-W., LEE K., TOUTANOVA K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (2019), NAACL '19, ACL, pp. 4171–4186. doi:10.18653/v1/N19-1423. 3, 16

[DMdO19] DIAS A. G., MILIOS E. E., DE OLIVEIRA M. C. F.: TRIVIR: A visualization system to support document retrieval with high recall. In *Proceedings of the ACM Symposium on Document Engineering* (2019), DocEng '19, ACM, pp. 1–10. doi:10.1145/3342558.3345401. 8, 17

[DSS*23] DATTA D., SELF N., SIMEONE J., MEADOWS A., OUTHWAITE W., WALKER L., ELMQVIST N., RAMAKRISHNAN N.: TimberSleuth: Visual anomaly detection with human feedback for mitigating the illegal timber trade. *Information Visualization* (Mar. 2023), 14738716231157081. doi:10.1177/14738716231157081. 16, 20

[EAKC*20] EL-ASSADY M., KEHLBECK R., COLLINS C., KEIM D. A., DEUSSEN O.: Semantic Concept Spaces: Guided topic model refinement using word-embedding projections. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (Jan. 2020), 1001–1011. doi:10.1109/TVCG.2019.2934654. 1, 8, 21, 22, 24

[EHA*22] ECKELT K., HINTERREITER A., ADELBERGER P., WALCHSHOFER C., DHANOA V., HUMER C., HECKMANN M., STEINPARZ C., STREIT M.: Visual exploration of relationships and structure in low-dimensional embeddings. *IEEE Transactions on Visualization and Computer Graphics* (2022). doi:10.1109/TVCG.2022.3156760. 5, 15

[EHR*14] ENDERT A., HOSSAIN M. S., RAMAKRISHNAN N., NORTH C., FIAUX P., ANDREWS C.: The human is the loop: New directions for visual analytics. *Journal of Intelligent Information Systems 43*, 3 (Dec. 2014), 411–435. doi:10.1007/s10844-014-0304-9. 1

[EMK*21] ESPADOTO M., MARTINS R. M., KERREN A., HIRATA N. S. T., TELEA A. C.: Toward a quantitative survey of dimension reduction techniques. *IEEE Transactions on Visualization and Computer Graphics 27*, 3 (Mar. 2021), 2153–2173. doi:10.1109/TVCG.2019.2944182. 2, 4

[ERT*17] ENDERT A., RIBARSKY W., TURKAY C., WONG B. W., NABNEY I., BLANCO I. D., ROSSI F.: The state of the art in integrating machine learning into visual analytics. *Computer Graphics Forum 36*, 8 (June 2017), 458–486. doi:10.1111/cgf.13092. 4, 22

[FCH*22] FENG Y., CHEN J., HUANG K., WONG J. K., YE H., ZHANG W., ZHU R., LUO X., CHEN W.: iPoet: Interactive painting poetry creation with visual multimodal analysis. *Journal of Visualization 25*, 3 (June 2022), 671–685. doi:10.1007/s12650-021-00780-0. 13, 14, 18, 20, 22

[FF20] FEKETE J.-D., FREIRE J.: Exploring reproducibility in visualization. *IEEE Computer Graphics and Applications 40*, 5 (Sept.–Oct. 2020), 108–119. doi:10.1109/MCG.2020.3006412. 25

[FFST19] FAVELIER G., FARAJ N., SUMMA B., TIERNY J.: Persistence Atlas for critical point variability in ensembles. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (Jan. 2019), 1152–1162. doi:10.1109/TVCG.2018.2864432. 9, 16

[FYC*22] FENG F., YANG Y., CER D., ARIVAZHAGAN N., WANG W.: Language-agnostic BERT sentence embedding. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (May 2022), ACL '22, ACL, pp. 878–891. doi:10.18653/v1/2022.acl-long.62. 1

[FZCM20] FUJIWARA T., ZHAO J., CHEN F., MA K.-L.: A visual analytics framework for contrastive network analysis. In *Proceedings of the 2020 IEEE Conference on Visual Analytics Science and Technology* (2020), VAST '20, IEEE, pp. 48–59. doi:10.1109/VAST50239.2020.00010. 7, 8, 9, 17, 18, 20

[GF18] GOYAL P., FERRARA E.: Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems 151* (July 2018), 78–94. doi:10.1016/j.knosys.2018.03.022. 1, 3, 4

[GFL*20] GUO R., FUJIWARA T., LI Y., LIMA K. M., SEN S., TRAN N. K., MA K.-L.: Comparative visual analytics for assessing medical records with sequence embedding. *Visual Informatics 4*, 2 (June 2020), 72–85. doi:10.1016/j.visinf.2020.04.001. 7, 10, 16

[GGW22] GROSSMANN N., GRÖLLER E., WALDNER M.: Concept splatters: Exploration of latent spaces based on human interpretable concepts. *Computers & Graphics 105* (June 2022), 73–84. doi:10.1016/j.cag.2022.04.013. 16, 17, 24

[GHM21] GHAI B., HOQUE M. N., MUELLER K.: WordBias: An interactive visual tool for discovering intersectional biases encoded in word embeddings. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (2021), CHI EA '21, ACM, pp. 1–7. doi:10.1145/3411763.3451587. 8, 15, 16, 20, 23

[GL16] GROVER A., LESKOVEC J.: node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016), KDD '16, ACM, pp. 855–864. doi:10.1145/2939672.2939754. 3

[Gro22] GROOTENDORST M.: BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022). doi:10.48550/arXiv.2203.05794. 20

[GWW19] GUO W., WANG J., WANG S.: Deep multimodal representation learning: A survey. *IEEE Access 7* (2019), 63373–63394. doi:10.1109/ACCESS.2019.2916887. 1, 10

[GZRP*22] GARCÍA-ZANABRIA G., RAIMUNDO M. M., POCO J., NERY M. B., SILVA C. T., ADORNO S., NONATO L. G.: CriPAV: Street-level crime patterns analysis and visualization. *IEEE Transactions on Visualization and Computer Graphics 28*, 12 (Dec. 2022), 4000–4015. doi:10.1109/TVCG.2021.3111146. 11, 12, 18, 20, 24

[HG18] HEIMERL F., GLEICHER M.: Interactive analysis of word vector embeddings. *Computer Graphics Forum 37*, 3 (June 2018), 253–265. doi:10.1111/cgf.13417. 8, 15, 18

[HGE22] HOQUE M. N., GHAI B., ELMQVIST N.: DramatVis Personae: Visual text analytics for identifying social biases in creative writing. In *Designing Interactive Systems Conference* (2022), DIS '22, ACM, pp. 1260–1276. doi:10.1145/3532106.3533526. 11

[HHS20] HOGRÄFER M., HEITZLER M., SCHULZ H.-J.: The state of the art in map-like visualization. *Computer Graphics Forum 39*, 3 (June 2020), 647–674. doi:10.1111/cgf.13672. 4

[HHS*23] HOQUE M. N., HE W., SHEKAR A. K., GOU L., REN L.: Visual Concept Programming: A visual analytics approach to injecting human intelligence at scale. *IEEE Transactions on Visualization and Computer Graphics 29*, 1 (Jan. 2023), 74–83. doi:10.1109/TVCG.2022.3209466. 16, 20

[HKD*21] HAGERER G., KIRCHHOFF M., DANNER H., PESCH R., GHOSH M., ROY A., ZHAO J., GROH G.: SocialVisTUM: An interactive visualization toolkit for correlated neural topic models on social media opinion mining. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing* (2021), RANLP '21, INCOMA Ltd., pp. 475–482. URL: https://aclanthology.org/2021.ranlp-1.54. 12, 15, 16, 18, 20

[HKMG22] HEIMERL F., KRALJ C., MOLLER T., GLEICHER M.: embComp: Visual interactive comparison of vector embeddings. *IEEE Transactions on Visualization and Computer Graphics 28*, 8 (Aug. 2022), 2953–2969. doi:10.1109/TVCG.2020.3045918. 14, 15

[HKPC18] HOHMAN F., KAHNG M., PIENTA R., CHAU D. H.: Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE Transactions on Visualization and Computer Graphics 25*, 8 (Aug. 2018), 2674–2693. doi:10.1109/TVCG.2018.2843369. 4

[HMW*15] HERZOG R., MEWES D., WAND M., GUIBAS L., SEIDEL H.-P.: LeSSS: Learned shared semantic spaces for relating multi-modal representations of 3D shapes. *Computer Graphics Forum 34*, 5 (Aug. 2015), 141–151. doi:10.1111/cgf.12703. 16, 18, 24

[HPvU*16] HÖLLT T., PEZZOTTI N., VAN UNEN V., KONING F., EISEMANN E., LELIEVELDT B. P., VILANOVA A.: Cytosplore: Interactive immune cell phenotyping for large single-cell datasets. *Computer Graphics Forum 35*, 3 (June 2016), 171–180. doi:10.1111/cgf.12893. 7, 9, 10, 18

[HPvU*18] HÖLLT T., PEZZOTTI N., VAN UNEN V., KONING F., LELIEVELDT B. P., VILANOVA A.: CyteGuide: Visual guidance for hierarchical single-cell analysis. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (Jan. 2018), 739–748. doi:10.1109/TVCG.2017.2744318. 10, 16, 18, 24

[HPX*21] HAN D., PAN J., XIE C., ZHAO X., LUO X., CHEN W.: A visual analytics approach for structural differences among graphs via deep learning. *IEEE Computer Graphics and Applications 41*, 5 (Sept.–Oct. 2021), 18–31. doi:10.1109/MCG.2021.3097799. 3, 18

[HPY00] HAN J., PEI J., YIN Y.: Mining frequent patterns without candidate generation. *ACM SIGMOD Record 29*, 2 (May 2000), 1–12. doi:10.1145/335191.335372. 20

[HRZ*20] HOU M., REN J., ZHANG D., KONG X., ZHANG D., XIA F.: Network embedding: Taxonomies, frameworks and applications. *Computer Science Review 38* (Nov. 2020), 100296. doi:10.1016/j.cosrev.2020.100296. 4

[HS97] HOCHREITER S., SCHMIDHUBER J.: Long short-term memory. *Neural Computation 9*, 8 (1997), 1735–1780. doi:10.1162/neco.1997.9.8.1735. 12

[HTAA*15] HUANG D., TORY M., ADRIEL ASENIERO B., BARTRAM L., BATEMAN S., CARPENDALE S., TANG A., WOODBURY R.: Personal visualization and personal visual analytics. *IEEE Transactions on Visualization and Computer Graphics 21*, 3 (Mar. 2015), 420–433. doi:10.1109/TVCG.2014.2359887. 15

[Hu18] HU Y.: Geo-text data and data-driven geospatial semantics. *Geography Compass 12*, 11 (2018), e12404. doi:10.1111/gec3.12404. 11

[HYL17] HAMILTON W. L., YING R., LESKOVEC J.: Representation learning on graphs: Methods and applications. *IEEE Data Engineering Bulletin 40*, 3 (Sept. 2017), 52–74. URL: http://sites.computer.org/debull/A17sept/issue1.htm. 1

[JCS*21] JIN Z., CHEN N., SHI Y., QIAN W., XU M., CAO N.: TrammelGraph: Visual graph abstraction for comparison. *Journal of Visualization 24*, 2 (Apr. 2021), 365–379. doi:10.1007/s12650-020-00706-2. 10, 24

[JCSM22] JASIM M., COLLINS C., SARVGHAD A., MAHYAR N.: Supporting serendipitous discovery and balanced analysis of online product reviews with interaction-driven metrics and bias-mitigating suggestions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2022), CHI '22, ACM, pp. 1–24. doi:10.1145/3491102.3517649. 6, 7, 12, 14, 15, 16

[JLB22] JEONG S., LIU S., BERGER M.: Interactively assessing disentanglement in GANs. *Computer Graphics Forum 41*, 3 (June 2022), 85–95. doi:10.1111/cgf.14524. 7

[JM20] JOSEPH K., MORGAN J.: When do word embeddings accurately reflect surveys on our beliefs about people? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (July 2020), ACL '20, ACL, pp. 4392–4415. doi:10.18653/v1/2020.acl-main.405. 23

[JSR*19] JI X., 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 25*, 6 (June 2019), 2181–2192. doi:10.1109/TVCG.2019.2903946. 1, 10, 22

[JWC*23] JIN Z., WANG X., CHENG F., SUN C., LIU Q., QU H.: ShortcutLens: A visual analytics approach for exploring shortcuts in natural language understanding dataset. *IEEE Transactions on Visualization and Computer Graphics* (2023). doi:10.1109/TVCG.2023.3236380. 8

[KAF*08] KEIM D. A., ANDRIENKO G., FEKETE J.-D., GÖRG C., KOHLHAMMER J., MELANÇON G.: Visual analytics: Definition, process, and challenges. In *Information Visualization: Human-Centered Issues and Perspectives*, vol. 4950 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2008, pp. 154–175. doi:10.1007/978-3-540-70956-5_7. 2

[KCD*19] KIM H., CHOI D., DRAKE B., ENDERT A., PARK H.: TopicSifter: Interactive search space reduction through targeted topic modeling. In *Proceedings of the 2019 IEEE Conference on Visual Analytics Science and Technology* (2019), VAST '19, IEEE, pp. 35–45. doi:10.1109/VAST47406.2019.8986922. 9, 10, 16

[KCK*19] 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 25*, 1 (Jan. 2019), 299–309. doi:10.1109/TVCG.2018.2865027. 7, 10, 14, 16, 20

[KFLCO13] KLEIMAN Y., FISH N., LANIR J., COHEN-OR D.: Dynamic maps for exploring and browsing shapes. *Computer Graphics Forum 32*, 5 (Aug. 2013), 187–196. doi:10.1111/cgf.12185. 13, 16

[KGM*22] KALAMARAS I., GLYKOS K., MEGALOOIKONOMOU V., VOTIS K., TZOVARAS D.: Graph-based visualization of sensitive medical data. *Multimedia Tools and Applications 81*, 1 (Jan. 2022), 209–236. doi:10.1007/s11042-021-10990-1. 10

[KK15] KUCHER K., KERREN A.: Text visualization techniques: Taxonomy, visual survey, and community insights. In *Proceedings of the IEEE Pacific Visualization Symposium* (2015), PacificVis '15, IEEE, pp. 117–121. doi:10.1109/PACIFICVIS.2015.7156366. 4, 5, 6

[KK19] KUCHER K., KERREN A.: Text visualization revisited: The state of the field in 2019. In *Poster Abstracts of the EG/VGTC Conference on Visualization* (2019), EuroVis '19, The Eurographics Association, pp. 29–31. doi:10.2312/eurp.20191138. 4

[KKLS17] KERREN A., KUCHER K., LI Y.-F., SCHREIBER F.: BioVis Explorer: A visual guide for biological data visualization techniques. *PLOS ONE 12*, 11 (Nov. 2017). doi:10.1371/journal.pone.0187341. 4

[KMK18] KUCHER K., MARTINS R. M., KERREN A.: Analysis of VINCI 2009–2017 proceedings. In *Proceedings of the 11th International Symposium on Visual Information Communication and Interaction* (2018), VINCI '18, ACM, pp. 97–101. doi:10.1145/3231622.3231641. 4

[KOK*18] KOYAMADA K., ONOUE Y., KIOKA M., UETSUJI T., BABA K.: Visualization of JOV abstracts. *Journal of Visualization 21*, 2 (Apr. 2018), 309–319. doi:10.1007/s12650-017-0451-5. 8, 15

[KØSV18] KUTUZOV A., ØVRELID L., SZYMANSKI T., VELLDAL E.: Diachronic word embeddings and semantic shifts: A survey. In *Proceedings of the International Conference on Computational Linguistics* (2018), COLING '18, ACL, pp. 1384–1397. URL: https://aclanthology.org/C18-1117. 1, 4

[KPK18] KUCHER K., PARADIS C., KERREN A.: The state of the art in sentiment visualization. *Computer Graphics Forum 37*, 1 (Feb. 2018), 71–96. doi:10.1111/cgf.13217. 4, 5, 6, 7

[KW18] KURZHALS K., WEISKOPF D.: Exploring the visualization design space with repertory grids. *Computer Graphics Forum 37*, 3 (June 2018), 133–144. doi:10.1111/cgf.13407. 16, 17, 20

[KW19] KINGMA D. P., WELLING M.: An introduction to variational autoencoders. *Foundations and Trends in Machine Learning 12*, 4 (2019), 307–392. doi:10.1561/2200000056. 11

[LBT*18] LIU S., BREMER P.-T., THIAGARAJAN J. J., SRIKUMAR V., WANG B., LIVNAT Y., PASCUCCI V.: Visual exploration of semantic relationships in neural word embeddings. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (Jan. 2018), 553–562. doi:10.1109/TVCG.2017.2745141. 8, 14, 20

[LCSEK19] LAUGHLIN B., COLLINS C., SANKARANARAYANAN K., EL-KHATIB K.: A visual analytics framework for adversarial text generation. In *Proceedings of the 2019 IEEE Symposium on Visualization for Cyber Security* (2019), VizSec '19. doi:10.1109/VizSec48167.2019.9161563. 8, 18

[LDL*20] LI C., DONG X., LIU W., SHENG S., QIAN A.: SSRD-Vis: Interactive visualization for event sequences summarization and rare detection. *Journal of Visualization 23*, 1 (Feb. 2020), 171–184. doi:10.1007/s12650-019-00609-x. 9, 18, 20, 24

[LGY*20] LI B., GOULD J., YANG Y., SARKIZOVA S., TABAKA M., ASHENBERG O., ROSEN Y., SLYPER M., KOWALCZYK M. S., VILLANI A.-C., TICKLE T., HACOHEN N., ROZENBLATT-ROSEN O., REGEV A.: Cumulus provides cloud-based data analysis for large-scale single-cell and single-nucleus RNA-seq. *Nature Methods 17*, 8 (Aug. 2020), 793–798. doi:10.1038/s41592-020-0905-x. 10

[LHZ*22] LIU Y., HU M., ZHANG R., XU T., WANG Y., ZHOU Z.: Visual aggregation of large multivariate networks with attribute-enhanced representation learning. *Neurocomputing 494* (July 2022), 320–335. doi:10.1016/j.neucom.2022.04.110. 16, 20

[LJLH19] LIU Y., JUN E., LI Q., HEER J.: Latent Space Cartography: Visual analysis of vector space embeddings. *Computer Graphics Forum 38*, 3 (June 2019), 67–78. doi:10.1111/cgf.13672. 16, 17

[LKB20] LIU Q., KUSNER M. J., BLUNSOM P.: A survey on contextual embeddings. *arXiv preprint arXiv:2003.07278* (2020). doi:10.48550/arXiv.2003.07278. 4, 7

[LKJ*20] 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 26*, 11 (Nov. 2020), 3133–3146. doi:10.1109/TVCG.2019.2922597. 11, 12, 17

[LM14] LE Q., MIKOLOV T.: Distributed representations of sentences and documents. In *Proceedings of the International Conference on Machine Learning* (2014), ICML '14, PMLR, pp. 1188–1196. URL: https://proceedings.mlr.press/v32/le14.html. 1, 3

[LMW*17] LIU S., MALJOVEC D., WANG B., BREMER P.-T., PASCUCCI V.: Visualizing high-dimensional data: Advances in the past decade. *IEEE Transactions on Visualization and Computer Graphics 23*, 3 (Mar. 2017), 1249–1268. doi:10.1109/TVCG.2016.2640960. 2, 4

[LNH*18] LI Q., NJOTOPRAWIRO K. S., HALEEM H., CHEN Q., YI C., MA X.: EmbeddingVis: A visual analytics approach to comparative network embedding inspection. In *Proceedings of the 2018 IEEE Conference on Visual Analytics Science and Technology* (2018), VAST '18, pp. 48–59. doi:10.1109/VAST.2018.8802454. 14, 16, 20

[LPH*20] LEKSCHAS F., PETERSON B., HAEHN D., MA E., GEHLENBORG N., PFISTER H.: PEAX: Interactive visual pattern search in sequential data using unsupervised deep representation learning. *Computer Graphics Forum 39*, 3 (June 2020), 167–179. doi:10.1111/cgf.13971. 10, 15, 24

[LSE21] LUO X., SCANDOLO L., EISEMANN E.: Texture Browser: Feature-based texture exploration. *Computer Graphics Forum 40*, 3 (June 2021), 99–109. doi:10.1111/cgf.14292. 13

[LTHL21] LIU H., TAO Y., HUANG W., LIN H.: Visual exploration of dependency graph in source code via embedding-based similarity. *Journal of Visualization 24*, 3 (June 2021), 565–581. doi:10.1007/s12650-020-00727-x. 9, 17, 20

[LWBM22] LIU Z., WANG Y., BERNARD J., MUNZNER T.: Visualizing graph neural networks with CorGIE: Corresponding a graph to its embedding. *IEEE Transactions on Visualization and Computer Graphics 28*, 6 (July 2022), 2500–2516. doi:10.1109/TVCG.2022.3148197. 8, 9, 15

[LWC*19] LIU S., WANG X., COLLINS C., DOU 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 25*, 7 (July 2019), 2482–2504. doi:10.1109/TVCG.2018.2834341. 4

[LWLZ17] LIU S., WANG X., LIU M., ZHU J.: Towards better analysis of machine learning models: A visual analytics perspective. *Visual Informatics 1*, 1 (Mar. 2017), 48–56. doi:10.1016/j.visinf.2017.01.006. 8

[LWZ*23] LI H., WANG J., ZHENG Y., WANG L., ZHANG W., SHEN H.-W.: Compressing and interpreting word embeddings with latent space regularization and interactive semantics probing. *Information Visualization 22*, 1 (Jan. 2023), 52–68. doi:10.1177/14738716221130338. 8, 15

[LXW*21] LI R., XIAO W., WANG L., JANG H., CARENINI G.: T3-Vis: Visual analytic for training and fine-tuning transformers in NLP. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (2021), EMNLP '21, ACL, pp. 220–230. doi:10.18653/v1/2021.emnlp-demo.26. 8, 17

[LY18] LI Y., YANG T.: Word embedding for understanding natural language: A survey. In *Guide to Big Data Applications*. Springer, 2018, pp. 83–104. doi:10.1007/978-3-319-53817-4_4. 4

[LYY*20] LI R., YIN C., YANG S., QIAN B., ZHANG P.: Marrying medical domain knowledge with deep learning on electronic health records: A deep visual analytics approach. *Journal of Medical Internet Research 22*, 9 (Sept. 2020), e20645. doi:10.2196/20645. 10, 16, 21

[LZ23] LI J., ZHOU C.-Q.: Incorporation of human knowledge into data embeddings to improve pattern significance and interpretability. *IEEE Transactions on Visualization and Computer Graphics 29*, 1 (Jan. 2023), 723–733. doi:10.1109/TVCG.2022.3209382. 17, 22

[MDL07] MAO Y., DILLON J., LEBANON G.: Sequential document visualization. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (Nov. 2007), 1208–1215. doi:10.1109/TVCG.2007.70592. 18, 20

[MHL*20] MIRANDA F., HOSSEINI M., LAGE M., DORAISWAMY H., DOVE G., SILVA C. T.: Urban Mosaic: Visual exploration of streetscapes using large-scale image data. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), CHI '20, ACM, pp. 1–15. doi:10.1145/3313831.3376399. 1, 11, 14, 20

[MLMRC18] MITHUN N. C., LI J., METZE F., ROY-CHOWDHURY A. K.: Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *Proceedings of the 2018 ACM International Conference on Multimedia Retrieval* (2018), ICMR '18, ACM, pp. 19–27. doi:10.1145/3206025.3206064. 1

[MM23] MAHMOOD S., MUELLER K.: Interactive subspace cluster analysis guided by semantic attribute associations. *IEEE Transactions on Visualization and Computer Graphics* (2023). doi:10.1109/TVCG.2023.3256376. 8, 14, 21

[MNS*23] MARTÍNEZ M. J., NAVEIRO R., SOTO A. J., TALAVANTE P., KIM LEE S.-H., GÓMEZ ARRAYAS R., FRANCO M., MAULEÓN P., LOZANO ORDÓÑEZ H., REVILLA LÓPEZ G., BERNABEI M., CAMPILLO N. E., PONZONI I.: Design of new dispersants using machine learning and visual analytics. *Polymers 15*, 5 (Jan. 2023), 1324. doi:10.3390/polym15051324. 13, 14

[MSC*13] MIKOLOV T., SUTSKEVER I., CHEN K., CORRADO G., DEAN J.: Distributed representations of words and phrases and their compositionality. In *Proceedings of the International Conference on Neural Information Processing Systems — Volume 2* (2013), NIPS '13, Curran Associates Inc., pp. 3111–3119. URL: https://dl.acm.org/doi/abs/10.5555/2999792.2999959. 1, 3, 11, 16

[Mül07] MÜLLER M.: Dynamic Time Warping. In *Information Retrieval for Music and Motion*, Müller M., (Ed.). Springer, 2007, pp. 69–84. doi:10.1007/978-3-540-74048-3_4. 11

[MWJ22] MEINECKE C., WRISLEY D. J., JÄNICKE S.: Explaining semi-supervised text alignment through visualization. *IEEE Transactions on Visualization and Computer Graphics 28*, 12 (Dec. 2022), 4797–4809. doi:10.1109/TVCG.2021.3105899. 8, 11, 14, 20, 21

[MWZ19] MOLINO P., WANG Y., ZHANG J.: Parallax: Visualizing and understanding the semantics of embedding spaces via algebraic formulae. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations* (2019), ACL '19, ACL, pp. 165–180. doi:10.18653/v1/P19-3028. 8, 15

[MZAD*20] MA C., ZHAO Y., AL-DOHUKI S., YANG J., YE X., KAMW F., AMIRUZZAMAN M.: GTMapLens: Interactive lens for geo-text data browsing on map. *Computer Graphics Forum 39*, 3 (June 2020), 469–481. doi:10.1111/cgf.13995. 11, 12, 16

[NA19] NONATO L. G., AUPETIT M.: Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE Transactions on Visualization and Computer Graphics 25*, 8 (Aug. 2019), 2650–2673. doi:10.1109/TVCG.2018.2846735. 4, 22

[NKWW22] NARECHANIA A., KARDUNI A., WESSLEN R., WALL E.: VITALITY: Promoting serendipitous discovery of academic literature with transformers & visual analytics. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (Jan. 2022), 486–496. doi:10.1109/TVCG.2021.3114820. 11, 16

[OMF*22] OUBENALI N., MESSAOUD S., FILIOT A., LAMER A., ANDREY P.: Visualization of medical concepts represented using word embeddings: A scoping review. *BMC Medical Informatics and Decision Making 22*, 1 (Dec. 2022), 83. doi:10.1186/s12911-022-01822-9. 10

[OMK21] OTTER D. W., MEDINA J. R., KALITA J. K.: A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems 32*, 2 (Feb. 2021), 604–624. doi:10.1109/TNNLS.2020.2979670. 3

[PARS14] PEROZZI B., AL-RFOU R., SKIENA S.: DeepWalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining* (2014), KDD '14, ACM, pp. 701–710. doi:10.1145/2623330.2623732. 12

[PC05] PIROLLI P., CARD S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of the International Conference on Intelligence Analysis* (2005), vol. 5, pp. 2–4. 4

[PCZ*21] PAN J., CHEN W., ZHAO X., ZHOU S., ZENG W., ZHU M., CHEN J., FU S., WU Y.: Exemplar-based layout fine-tuning for node-link diagrams. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (Feb. 2021), 1655–1665. doi:10.1109/TVCG.2020.3030393. 9, 10, 15, 16, 17, 22

[PDD*22] PARK H., DAS N., DUGGAL R., WRIGHT A. P., SHAIKH O., HOHMAN F., POLO CHAU D. H.: NeuroCartography: Scalable automatic visual summarization of concepts in deep neural networks. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (Jan. 2022), 813–823. doi:10.1109/TVCG.2021.3114858. 6, 8, 9, 16, 17, 18, 24

[PdSP*22] POCCO X., DA SILVA T., POCO J., NONATO L. G., GOMEZ-NIETO E.: Exploring scientific literature by textual and image content using DRIFT. *Computers & Graphics 103* (Apr. 2022), 140–152. doi:10.1016/j.cag.2022.02.005. 5, 8, 16, 18, 20

[PKL*18] 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 24*, 1 (Jan. 2018), 361–370. doi:10.1109/TVCG.2017.2744478. 8, 17, 22

[PPV*21] POCCO X., POCO J., VIANA M., DE PAULA R., NONATO L. G., GOMEZ-NIETO E.: DRIFT: A visual analytic tool for scientific literature exploration based on textual and image content. In *Proceedings of the 34th SIBGRAPI Conference on Graphics, Patterns and Images* (2021), SIBGRAPI '21, IEEE, pp. 136–143. doi:10.1109/SIBGRAPI54419.2021.00027. 5

[PS21] PRZYBYŁA P., SOTO A. J.: When classification accuracy is not enough: Explaining news credibility assessment. *Information Processing & Management 58*, 5 (Sept. 2021), 102653. doi:10.1016/j.ipm.2021.102653. 11

[PSM14] PENNINGTON J., SOCHER R., MANNING C. D.: GloVe: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)* (2014), pp. 1532–1543. URL: http://www.aclweb.org/anthology/D14-1162. 16

[QLL*22] QU D., LV C., LIN Y., ZHANG H., WANG R.: AirLens: Multi-level visual exploration of air quality evolution in urban agglomerations. *Computer Graphics Forum 41*, 3 (June 2022), 223–234. doi:10.1111/cgf.14535. 13, 16

[RCPW21] RATHORE A., CHALAPATHI N., PALANDE S., WANG B.: TopoAct: Visually Exploring the Shape of Activations in Deep Learning. *Computer Graphics Forum 40*, 1 (2021), 382–397. doi:10.1111/cgf.14195. 9

[RMH*22] RULFF J., MIRANDA F., HOSSEINI M., LAGE M., CARTWRIGHT M., DOVE G., BELLO J., SILVA C. T.: Urban Rhapsody: Large-scale exploration of urban soundscapes. *Computer Graphics Forum 41*, 3 (June 2022), 209–221. doi:10.1111/cgf.14534. 6, 11, 16, 18, 20, 24

[Rob07] ROBERTS J. C.: State of the art: Coordinated & multiple views in exploratory visualization. In *Proceedings of the Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization* (2007), CMV 2007, IEEE, pp. 61–71. doi:10.1109/CMV.2007.20. 2

[RPSM22] REZAEIPOURFARSANGI S., PEI N., SHERKAT E., MILIOS E.: Interactive clustering and high-recall information retrieval using language models. In *Proceedings of the 2022 International Conference on Advanced Visual Interfaces* (2022), AVI '22, ACM, pp. 1–5. doi:10.1145/3531073.3531174. 17

[RRM*21] RENOUST B., REN H., MELANÇON G., VIAUD M.-L., SATOH S.: A multimedia document browser based on multilayer networks. *Multimedia Tools and Applications 80*, 15 (June 2021), 22551–22588. doi:10.1007/s11042-020-09872-9. 14

[RSBV21] RAMAN N., SHAH S., BALCH T., VELOSO M.: ViziTex: Interactive visual sense-making of text corpora. In *Proceedings of the Second Workshop on Data Science with Human in the Loop: Language Advances* (2021), DaSH '21, ACL, pp. 16–23. doi:10.18653/v1/2021.dash-1.3. 8, 15

[RSL*22] RISSAKI A., SCARONE B., LIU D., PANDEY A., KLEIN B., ELIASSI-RAD T., BORKIN M. A.: BiaScope: Visual unfairness diagnosis for graph embeddings. In *Proceedings of the 2022 IEEE Symposium on Visualization in Data Science* (2022), VDS '22, IEEE, pp. 27–36. doi:10.1109/VDS57266.2022.00008. 6, 8, 14, 15, 16, 17, 20

[Rus16] RUSSELL D. M.: Simple is good: Observations of visualization use amongst the Big Data digerati. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* (2016), AVI '16, ACM, pp. 7–12. doi:10.1145/2909132.2933287. 14

[RvdHD*15] RAIDOU R., VAN DER HEIDE U., DINH C., GHOBADI G., KALLEHAUGE J., BREEUWER M., VILANOVA A.: Visual analytics for the exploration of tumor tissue characterization. *Computer Graphics Forum 34*, 3 (July 2015), 11–20. doi:10.1111/cgf.12613. 5

[SCR*23] SEVASTJANOVA R., CAKMAK E., RAVFOGEL S., COTTERELL R., EL-ASSADY M.: Visual comparison of language model adaptation. *IEEE Transactions on Visualization and Computer Graphics 29*, 1 (Jan. 2023), 1178–1188. doi:10.1109/TVCG.2022.3209458. 8, 14, 17, 21

[SDXR22] SONG H., DAI Z., XU P., REN L.: Interactive visual pattern search on graph data via graph representation learning. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (Jan. 2022), 335–345. doi:10.1109/TVCG.2021.3114857. 9, 10, 15, 16

[SEAG*21] SPERRLE F., EL-ASSADY M., GUO G., BORGO R., CHAU D. H., ENDERT A., KEIM D. A.: A survey of human-centered evaluations in human-centered machine learning. *Computer Graphics Forum 40*, 3 (June 2021), 543–568. doi:10.1111/cgf.14329. 4

[SGL22] SOHNS J.-T., GARTH C., LEITTE H.: Decision Boundary Visualization for Counterfactual Reasoning. *Computer Graphics Forum* (2022). doi:10.1111/cgf.14650. 5, 8, 17

[SGR*20] SUH J., GHORASHI S., RAMOS G., CHEN N.-C., DRUCKER S., VERWEY J., SIMARD P.: AnchorViz: Facilitating semantic data exploration and concept discovery for interactive machine learning. *ACM Transactions on Interactive Intelligent Systems 10*, 1 (Mar. 2020), 1–38. doi:10.1145/3241379. 7, 8, 18

[SH20] SHRIVASTAVA A., HEER J.: iSeqL: Interactive sequence learning. In *Proceedings of the 25th International Conference on Intelligent User Interfaces* (2020), IUI '20, ACM, pp. 43–54. doi:10.1145/3377325.3377503. 8, 18, 23

[SJB*17] STOFFEL F., JENTNER W., BEHRISCH M., FUCHS J., KEIM D.: Interactive ambiguity resolution of named entities in fictional literature. *Computer Graphics Forum 36*, 3 (June 2017), 189–200. doi:10.1111/cgf.13179. 8, 17, 21

[SKB*21] SEVASTJANOVA R., KALOULI A.-L., BECK C., SCHÄFER H., EL-ASSADY M.: Explaining contextualization in language models using visual analytics. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (2021), ACL-IJCNLP '21, ACL, pp. 464–476. doi:10.18653/v1/2021.acl-long.39. 8, 20

[SKB*22] SEVASTJANOVA R., KALOULI A., BECK C., HAUPTMANN H., EL-ASSADY M.: LMFingerprints: Visual explanations of language model embedding spaces through layerwise contextualization scores. *Computer Graphics Forum 41*, 3 (June 2022), 295–307. doi:10.1111/cgf.14541. 8, 13, 18, 23

[SKKC19] SACHA D., KRAUS M., KEIM D. A., CHEN M.: VIS4ML: An ontology for visual analytics assisted machine learning. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2019), 385–395. doi:10.1109/TVCG.2018.2864838. 7, 17

[SLN*19] SHOEMARK P., LIZA F. F., NGUYEN D., HALE S., MCGILLIVRAY B.: Room to Glo: A systematic comparison of semantic change detection approaches with word embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (2019), EMNLP-IJCNLP '19, ACL, pp. 66–76. doi:10.18653/v1/D19-1007. 24

[SNHS13] SCHULZ H.-J., NOCKE T., HEITZLER M., SCHUMANN H.: A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (Dec. 2013), 2366–2375. doi:10.1109/TVCG.2013.120. 17

[SNP*22] STYVE L., NAVARRA C., PETERSEN J. M., NESET T.-S., VROTSOU K.: A visual analytics pipeline for the identification and exploration of extreme weather events from social media data. *Climate 10*, 11 (Nov. 2022), 174. doi:10.3390/cli10110174. 11

[SSJ*22] SOHNS J.-T., SCHMITT M., JIRASEK F., HASSE H., LEITTE H.: Attribute-based explanation of non-linear embeddings of high-dimensional data. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (Jan. 2022), 540–550. doi:10.1109/TVCG.2021.3114870. 8, 17, 23

[SSKEA21] SPERRLE F., SCHÄFER H., KEIM D., EL-ASSADY M.: Learning contextualized user preferences for co-adaptive guidance in mixed-initiative topic model refinement. *Computer Graphics Forum 40*, 3 (June 2021), 215–226. doi:10.1111/cgf.14301. 8, 20

[SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G., KEIM D. A.: Knowledge generation model for visual analytics. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (Dec. 2014), 1604–1613. doi:10.1109/TVCG.2014.2346481. 2

[STN*16] SMILKOV D., THORAT N., NICHOLSON C., REIF E., VIÉGAS F. B., WATTENBERG M.: Embedding Projector: Interactive visualization and interpretation of embeddings. In *Proceedings of the NIPS 2016 Workshop on Interpretable Machine Learning in Complex Systems* (November 2016), arXiv. arXiv:1611.05469 [cs, stat]. URL: http://arxiv.org/abs/1611.05469. 14, 16, 17

[Str35] STROOP J. R.: Studies of interference in serial verbal reactions. *Journal of Experimental Psychology 18* (1935), 643–662. doi:10.1037/h0054651. 22

[SWP22] SIVARAMAN V., WU Y., PERER A.: Emblaze: Illuminating machine learning representations through interactive comparison of embedding spaces. In *Proceedings of the 27th International Conference on Intelligent User Interfaces* (2022), IUI '22, ACM, pp. 418–432. doi:10.1145/3490099.3511137. 13, 14, 16, 18, 23

[SXG*19] SUN Z., XING W., GUO W., KIM S., LI H., LI W., WU J., ZHANG Y., CHENG B., CHENG S.: A survey on dimension reduction algorithms in big data visualization. In *Cloud Computing, Smart Grid and Innovative Frontiers in Telecommunications*. Springer, 2019, pp. 375–395. doi:10.1007/978-3-030-48513-9_31. 4

[SZS*17] SACHA D., ZHANG L., SEDLMAIR M., LEE J. A., PELTONEN J., WEISKOPF D., NORTH S. C., KEIM D. A.: Visual interaction with dimensionality reduction: A structured literature analysis. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (Jan. 2017), 241–250. doi:10.1109/TVCG.2016.2598495. 2, 4, 6, 9

[TCS*22] TKACHEV G., CUTURA R., SEDLMAIR M., FREY S., ERTL T.: Metaphorical Visualization: Mapping data to familiar concepts. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts* (2022), CHI EA '22, ACM, pp. 1–10. doi:10.1145/3491101.3516393. 10, 15

[TWB*20] TENNEY I., WEXLER J., BASTINGS J., BOLUKBASI T., COENEN A., GEHRMANN S., JIANG E., PUSHKARNA M., RADEBAUGH C., REIF E., YUAN A.: The Language Interpretability Tool: Extensible, interactive visualizations and analysis for NLP models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (2020), EMNLP '20, ACL, pp. 107–118. doi:10.18653/v1/2020.emnlp-demos.15. 8, 17

[VLL*10] VINCENT P., LAROCHELLE H., LAJOIE I., BENGIO Y., MANZAGOL P.-A., BOTTOU L.: Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research 11*, 12 (2010). URL: https://jmlr.org/beta/papers/v11/vincent10a.html. 12

[VMZL22] VOIGT H., MEUSCHKE M., ZARRIESS S., LAWONN K.: KeywordScape: Visual document exploration using contextualized keyword embeddings. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (2022), ACL, pp. 137–147. URL: https://aclanthology.org/2022.emnlp-demos.14. 8, 18

[vUHP*17] VAN UNEN V., HÖLLT T., PEZZOTTI N., LI N., REINDERS M. J., EISEMANN E., KONING F., VILANOVA A., LELIEVELDT B. P.: Visual analysis of mass cytometry data by hierarchical stochastic neighbour embedding reveals rare cell types. *Nature Communications 8* (2017), 1740. doi:10.1038/s41467-017-01689-9. 9

[WHC*23] WANG Q., HUANG K., CHANDAK P., ZITNIK M., GEHLENBORG N.: Extending the nested model for user-centric XAI: A design study on GNN-based drug repurposing. *IEEE Transactions on Visualization and Computer Graphics 29*, 1 (Jan. 2023), 1266–1276. doi:10.1109/TVCG.2022.3209435. 14, 20

[WJM*22] WITSCHARD D., JUSUFI I., MARTINS R. M., KUCHER K., KERREN A.: Interactive optimization of embedding-based text similarity calculations. *Information Visualization 21*, 4 (Oct. 2022), 335–353. doi:10.1177/14738716221114372. 3, 9, 13, 14, 16, 17, 20, 24

[WJMK21] WITSCHARD D., JUSUFI I., MARTINS R. M., KERREN A.: A statement report on the use of multiple embeddings for visual analytics of multivariate networks. In *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP '21) — Volume 3: IVAPP* (2021), IVAPP '21, INSTICC, SciTePress, pp. 219–223. doi:10.5220/0010314602190223. 4

[WKM18] WENDLANDT L., KUMMERFELD J. K., MIHALCEA R.: Factors influencing the surprising instability of word embeddings. In *Proceedings of the 2018 Conference of the North American Chapter of the*

Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers) (2018), NAACL '18, ACL, pp. 2092–2102. doi:10.18653/v1/N18-1190. 23

[WLHO19] WANG Y., LAW N., HEMBERG E., O'REILLY U.-M.: Using detailed access trajectories for learning behavior analysis. In *Proceedings of the 9th International Conference on Learning Analytics & Knowledge* (2019), LAK '19, ACM, pp. 290–299. doi:10.1145/3303772.3303781. 11, 20

[WMW*22] WANG B., MA K., WU L., QIU Q., XIE Z., TAO L.: Visual analytics and information extraction of geological content for text-based mineral exploration reports. *Ore Geology Reviews 144* (May 2022), 104818. doi:10.1016/j.oregeorev.2022.104818. 13, 18

[WMWG17] WANG Q., MAO Z., WANG B., GUO L.: Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering 29*, 12 (Dec. 2017), 2724–2743. doi:10.1109/TKDE.2017.2754499. 1, 4

[WNT*20] WU H.-Y., NIEDERMANN B., TAKAHASHI S., ROBERTS M. J., NÖLLENBURG M.: A survey on transit map layout–From design, machine, and human perspectives. *Computer Graphics Forum 39*, 3 (June 2020), 619–646. doi:10.1111/cgf.13672. 4

[Woh14] WOHLIN C.: Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the International Conference on Evaluation and Assessment in Software Engineering* (2014), EASE '14, ACM, pp. 38:1–38:10. doi:10.1145/2601248.2601268. 5

[WSP*21] WRIGHT A. P., SHAIKH O., PARK H., EPPERSON W., AHMED M., PINEL S., CHAU D. H. P., YANG D.: RECAST: Enabling user recourse and interpretability of toxicity detection models with interactive visualization. *Proceedings of the ACM on Human-Computer Interaction 5*, CSCW1 (Apr. 2021), 1–26. doi:10.1145/3449280. 11, 15

[WWC*21] WANG J., WU J., CAO A., ZHOU Z., ZHANG H., WU Y.: Tac-Miner: Visual tactic mining for multiple table tennis matches. *IEEE Transactions on Visualization and Computer Graphics 27*, 6 (June 2021), 2770–2782. doi:10.1109/TVCG.2021.3074576. 13, 14, 20, 22

[WWS*21] WU A., WANG Y., SHU X., MORITZ D., CUI W., ZHANG H., ZHANG D., QU H.: AI4VIS: Survey on artificial intelligence approaches for data visualization. *IEEE Transactions on Visualization and Computer Graphics* (2021). doi:10.1109/TVCG.2021.3099002. 4

[WZY20] WANG J., ZHANG W., YANG H.: SCANViz: Interpreting the Symbol-Concept Association Captured by Deep Neural Networks through Visual Analytics. In *Proceedings of the 2020 IEEE Pacific Visualization Symposium* (2020), PacificVis '20, IEEE, pp. 51–60. doi:10.1109/PacificVis48177.2020.3542. 9

[XBL*18] XU S., BRYAN C., LI J. K., ZHAO J., MA K.-L.: Chart Constellations: Effective chart summarization for collaborative and multi-user analyses. *Computer Graphics Forum 37*, 3 (June 2018), 75–86. doi:10.1111/cgf.13402. 7, 12, 20, 22

[XO21] XIAO T., ONOUE Y.: Visualization of topic transitions in SNSs using document embedding and dimensionality reduction. In *Proceedings of the 2021 IEEE 14th Pacific Visualization Symposium* (2021), PacificVis '21, IEEE, pp. 216–220. doi:10.1109/PacificVis52677.2021.00035. 11

[XTL*21] XIE P., TAO W., LI J., HUANG W., CHEN S.: Exploring multi-dimensional data via subset embedding. *Computer Graphics Forum 40*, 3 (June 2021), 75–86. doi:10.1111/cgf.14290. 10, 14, 17

[XTYL18] XU J., TAO Y., YAN Y., LIN H.: VAUT: A visual analytics system of spatiotemporal urban topics in reviews. *Journal of Visualization 21*, 3 (June 2018), 471–484. doi:10.1007/s12650-017-0464-0. 11, 12, 16, 20

[XTYL20] XU J., TAO Y., YAN Y., LIN H.: Exploring evolution of dynamic networks via diachronic node embeddings. *IEEE Transactions on Visualization and Computer Graphics 26*, 7 (July 2020), 2387–2402. doi:10.1109/TVCG.2018.2887230. 9, 16, 17, 20, 23, 24

[XWX*22] XIE L., WU Z., XU P., LI W., MA X., LI Q.: RoleSeer: Understanding informal social role changes in MMORPGs via visual analytics. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2022), CHI '22, ACM, pp. 1–17. doi:10.1145/3491102.3517712. 11, 13, 14, 16, 20, 23

[XXM19] XIE C., XU W., MUELLER K.: A visual analytics framework for the detection of anomalous call stack trees in high performance computing applications. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (Jan. 2019), 215–224. doi:10.1109/TVCG.2018.2865026. 9, 16, 20

[XZL*22] XUE G., ZHONG M., LI J., CHEN J., ZHAI C., KONG R.: Dynamic network embedding survey. *Neurocomputing 472* (Feb. 2022), 212–223. doi:10.1016/j.neucom.2021.03.138. 7

[YHZ22] YE Y., HUANG R., ZENG W.: VISAtlas: An image-based exploration and query system for large visualization collections via neural image embedding. *IEEE Transactions on Visualization and Computer Graphics* (2022). doi:10.1109/TVCG.2022.3229023. 3, 10, 16

[ZDL21] ZHANG T., DUAN X., LI Y.: Unveiling transit mobility structure towards sustainable cities: An integrated graph embedding approach. *Sustainable Cities and Society 72* (Sept. 2021), 103027. doi:10.1016/j.scs.2021.103027. 12, 18, 20

[ZFC*21] ZHANG X., FUJIWARA T., CHANDRASEGARAN S., BRUNDAGE M. P., SEXTON T., DIMA A., MA K.-L.: A visual analytics approach for the diagnosis of heterogeneous and multidimensional machine maintenance data. In *Proceedings of the 2021 IEEE 14th Pacific Visualization Symposium* (2021), PacificVis '21, IEEE, pp. 196–205. doi:10.1109/PacificVis52677.2021.00033. 7, 9, 12

[ZJQH22] ZHOU Y., JIANG R., QIN H., HU H.: Representation and analysis of time-series data via deep embedding and visual exploration. *Journal of Visualization* (Oct. 2022). doi:10.1007/s12650-022-00890-3. 15, 16, 20

[ZSHL18] ZHOU Z., SHI C., HU M., LIU Y.: Visual ranking of academic influence via paper citation. *Journal of Visual Languages & Computing 48* (Oct. 2018), 134–143. doi:10.1016/j.jvlc.2018.08.007. 8, 18

[ZWW*22] ZENG H., WANG X., WANG Y., WU A., PONG T.-C., QU H.: GestureLens: Visual analysis of gestures in presentation videos. *IEEE Transactions on Visualization and Computer Graphics* (2022). doi:10.1109/TVCG.2022.3169175. 11, 20

[ZYHD20] ZHANG C., YANG Z., HE X., DENG L.: Multimodal intelligence: Representation learning, information fusion, and applications. *IEEE Journal of Selected Topics in Signal Processing 14*, 3 (Mar. 2020), 478–493. doi:10.1109/JSTSP.2020.2987728. 24

[ZYZZ20] ZHANG D., YIN J., ZHU X., ZHANG C.: Network representation learning: A survey. *IEEE Transactions on Big Data 6*, 1 (Mar. 2020), 3–28. doi:10.1109/TBDATA.2018.2850013. 1

[ZZL*22] ZENG X., ZHOU H., LI Z., ZHANG C., LIN J., XIA J., YANG Y., KUI X.: iHELP: Interactive hierarchical linear projections for interpreting non-linear projections. *Journal of Visualization* (Nov. 2022). doi:10.1007/s12650-022-00900-4. 5, 17