



REGULAR PAPER

Wei Zeng · Ao Dong · Xi Chen · Zhang-lin Cheng

VIStory: interactive storyboard for exploring visual information in scientific publications

Received: 6 February 2020 / Revised: 20 May 2020 / Accepted: 1 July 2020 / Published online: 16 August 2020
© The Visualization Society of Japan 2020

Abstract Many visual analytics have been developed for examining scientific publications comprising wealthy data such as authors and citations. The studies provide unprecedented insights on a variety of applications, e.g., literature review and collaboration analysis. However, visual information (e.g., figures) that is widely employed for storytelling and methods description are often neglected. We present *VIStory*, an interactive storyboard for exploring visual information in scientific publications. We harvest a new dataset of a large corpora of figures, using an automatic figure extraction method. Each figure contains various attributes such as dominant color and width/height ratio, together with faceted metadata of the publication including venues, authors, and keywords. To depict these information, we develop an intuitive interface consisting of three components: (1) Faceted View enables efficient query by publication metadata, benefiting from a nested table structure, (2) Storyboard View arranges paper rings—a well-designed glyph for depicting figure attributes, in a themeriver layout to reveal temporal trends, and (3) Endgame View presents a highlighted figure together with the publication metadata. We illustrate the applicability of *VIStory* with case studies on two datasets, i.e., 10-year IEEE VIS publications, and publications by a research team at CVPR, ICCV, and ECCV conferences. Quantitative and qualitative results from a formal user study demonstrate the efficiency of *VIStory* in exploring visual information in scientific publications.

Keywords Document visualization · Image browser · Faceted metadata

1 Introduction

Publications are one of the most important outcomes of scientific research. Together with the development of science itself, substantial amounts of scientific publications have been generated. Though digital libraries like *Google Scholar* and *Microsoft Academic* provide powerful searching and browsing functionalities, they are often found ineffective for high-level tasks such as collaboration analysis. Visual analytics has gained intense interest in exploring scientific publications, as it can enable human cognition and reasoning with machine's powerful computing capacity (Keim et al. 2008). Vast amounts of visual analytics have been developed that facilitate applications including literature review and citation analysis (e.g., Chou and Yang 2011; Heimerl et al. 2016; Wu et al. 2015; Latif and Beck 2018).

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s12650-020-00688-1>) contains supplementary material, which is available to authorized users.

W. Zeng (✉) · A. Dong · X. Chen · Z. Cheng
Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China
E-mail: wei.zeng@siat.ac.cn

A. Dong · X. Chen
University of Chinese Academy of Sciences, Beijing, China

Visual information (e.g., figures) are typically employed in scientific visualizations for describing facts, methods, or telling stories (Strobelt et al. 2009). Specifically, in fields such as visualization, research processes generate imagery data that can substantially reflect content and quality of the research (Chen et al. 2009). Taking visualization publications for an example, studying the visual information can benefit the field from multiple perspectives: to generate compact visual representations (e.g., Strobelt et al. 2009), to guide design processes to make memorable, recognizable, and recallable visualizations (e.g., Borkin et al. 2013, 2016), and to provide quality metrics for evaluating visualizations (e.g., Jänicke and Chen 2010; Matzen et al. 2018).

Many surveys on subfields of visualization, e.g., *Treevis.net* for tree visualization (Schulz 2011), *TimeViz Browser* (Tominski and Aigner 2015) for time-series visualization, *Text Visualization Browser* (Kucher and Kerren 2015) for text visualization, utilize imagery thumbnails to depict visualization techniques. The visual information enable users to quickly get an overview of the field, facilitate teaching, and find related work based on various categories defined in a survey taxonomy (Kucher and Kerren 2015). However, existing visual analytics for scientific publications typically focus on metadata of scientific publications such as authors and citations (e.g., Beck et al. 2016; Federico et al. 2017), while neglect the visual information. There is an emerging need for a survey system that enables quick and comprehensive exploration of visual information in scientific publications.

We develop an interactive storyboard, namely *VIStory*, that fulfills the requirement. To ease the maintenance pressure for such a visual survey system, we develop an automatic method to extract figures and multi-faceted metadata (including authors, keywords, venues, etc.) from scientific publications, and we construct a nested table structure to support efficient query of the visual information (Sect. 4). Next, we design an intuitive visual interface that adopt well-established visualization techniques of glyph design for visual information, themeriver layout for temporal variation, faceted exploration for multi-faceted analysis, and endgame view for details exploration (Sect. 5). We elaborate the utility and effectiveness of *VIStory* via case studies on ten-year IEEE VIS publications from 2009–2018, and publications by a multimedia group at three major computer vision conferences, i.e., CVPR, ICCV, and ECCV (Sect. 6). The studies reveal some interesting patterns, such as the emerging of machine learning topics in visualization conferences. We also conduct a formal user study in comparison with a state-of-the-art visual survey system—*SurVis* (Beck et al. 2016). The quantitative results demonstrate the efficiency, and the qualitative feedback shows the preference of *VIStory* (Sect. 7).

In summary, we contribute to the visual exploration of scientific publications in the following way:

- First, we propose a new perspective of exploring scientific publications, i.e., analyzing the visual information. Specifically, we curate a total of 11,568 figures from ten-year IEEE VIS publications in 2009–2018, using an automatic figure and metadata extraction method. We will release the dataset to foster future research.
- Second, we develop an interactive storyboard—*VIStory* that facilitates the exploration of visual information in scientific publications. *VIStory* integrates a compact glyph design of paper ring to represent multi-dimensional attributes of figures in one publication. The paper rings are arranged in a themeriver layout for depicting temporal trends. Faceted views and endgame views are also incorporated to support multi-faceted analysis and details-on-demand exploration.
- Lastly, we present three case studies conducted on the collected figures, to support real-world usage scenarios: author profile probe, VIS trend analysis, and research impact exploration. We figured out some interesting patterns, such as emerging topic of Machine Learning in IEEE VIS publications, which help users quickly understand a research field.

2 Related work

We group related work in the three categories: *Visual document analysis* (discussing studies of developing visual analytics for documentation analysis); *Image browser* (discussing general methods for visually exploring images); *Visualization taxonomies* (discussing recent trends of understanding visualization field by analyzing visualization publications).

2.1 Visual document analysis

The plethora of scientific publications poses challenges for the literature review. Though digital libraries such as *Google Scholar* and *Microsoft Academic* enable search by concepts or keywords, researchers can easily lose focus as little abstraction of tremendous raw publications are made by the digital libraries (Chou and Yang 2011). Many visual analytics have been developed to fill the gap. The systems can be categorized by data types of multi-faceted metadata including authors, references, title, and publication date (Federico et al. 2017). Exemplary work include Jigsaw (Görg et al. 2013) and HierarchicalTopics (Dou et al. 2013) for text analysis, PaperVis (Chou and Yang 2011) and CiteRivers (Heimerl et al. 2016) for citation, and egoSlider (Wu et al. 2015) and Vis Author Profile (Latif and Beck 2018) for authorship analysis. Coupled with advanced analysis techniques and intuitive visual designs, these systems have been proven effective in facilitating the understanding and assessment of scientific publications (Federico et al. 2017).

However, only a few visualizations are developed for depicting visual information in scientific publications. Strobel et al. (2009) organized key figures and important terms in a compact manner to generate an abstraction for documents. Schulz (2011) collected all tree visualization techniques, and developed a reference system that supports interactive exploration. Chen et al. (2020) analyzed composition and configuration patterns in multiple-view visualizations collected from IEEE VIS, EuroVis, and PacificVis conferences. A similar reference website was later developed for text visualization (Kucher and Kerren 2015). Unfortunately, the visualizations are either suitable for only a small amount of documents (Strobel et al. 2009), or relying heavily on developer expertises for maintenance (Schulz 2011; Kucher and Kerren 2015). Instead, this work aims for an interactive storyboard for vast amounts of figures automatically extracted from scientific publications.

2.2 Image browser

To visualize massive amounts of figures calls for effective image browser. A common approach is to organize images in a layout based on pairwise image similarities (Plant and Schaefer 2011). The layout has many variations, such as Neighbor-Joining tree (Eler et al. 2009), multi-dimensional scaling (Joia et al. 2011), Voronoi treemap (Tan et al. 2012), or picture collage (Liang et al. 2018). Some image browsers make use of images' semantic information that can be generated from conventional image annotation (Yang et al. 2006), emerging deep learning (Xie et al. 2018), and mutual information (Zeng et al. 2019). Besides similarities and semantics, images can also comprise multi-dimensional metadata such as place and categories, which can be used to facilitate searching and browsing (Corput and Wijk 2016). MediaTable (Rooij et al. 2010) arranged all images and associate metadata in a tabular layout. PICTuReVis (van der Corput and van Wijk 2017) showed that relations among people can be revealed based on image collections. StreetVizor (Shen et al. 2018) compared attributes extracted from spatial-dependent street views in cities.

This work adopts a conventional approach of browsing images using multi-faceted metadata (Yee et al. 2003), which is naturally compatible with conventional visualization mantra 'overview first, zoom and filter, then details on demand' (Shneiderman 1996). Moreover, publication figures exhibit multi-dimensional and temporal properties, requiring new visual design for depicting high-level patterns such as temporal variations. We employ glyph-based designs to intuitively depict multi-dimensional attributes, and arrange them in a themeriver layout to reveal temporal variations.

2.3 Visualization taxonomies

This work can facilitate the understanding of visualization field, by studying past 10-year visualization publications in IEEE VIS. We follow the suggestions from (Isenberg et al. 2017b) to study visualization publications by taxonomy. In addition to traditional taxonomy factors such as venues and topics, visualization researchers have identified more reasonable factors including data type (Shneiderman 1996), design models (Tory and Möller 2004), and data encodings (Rodrigues et al. 2006). Borkin et al. (2013) adopted a more detailed visualization taxonomy that categorize statistical charts based on visual encodings and perceptual tasks. Recently, Isenberg et al. (2017b) examined the keywords in IEEE VIS publications. They identified shortcomings of existing visualization taxonomy and proposed common terminologies. An online query tool (<http://keyvis.org/>) was also released to the public. We follow the work to categorize visualization publications based on keywords, e.g., volume rendering and user study. Furthermore, we develop an interactive storyboard to provide intuitive visual exploration of IEEE VIS publications.

3 Requirements and system overview

This work aims to explore visual information in scientific publications, which can benefit different users in addressing various tasks, e.g.,

- *Students* would like to compare profiles of different authors/topics when choosing supervisors/research topics.
- *Researchers* wish to understand what the others have developed when expanding their research areas.
- *Reviewers* may need to confirm if a visual design has been published when reviewing papers.

3.1 Requirements

We opt to develop an interactive storyboard to facilitate exploration, which shall meet the following requirements:

- *R1. Automation*: The system should enable automatic collection of visual information from scientific publications. In this way, minimum maintenance efforts are required, rather than heavy manual work by professional experts.
- *R2. Multi-faceted analysis*: The system should support multi-faceted analysis to meet different tasks, e.g., to help students find active researchers in the field, or to assist researchers in revealing trend of research topics in recent years.
- *R3. Intuitive visual design*: The system should incorporate intuitive visual design to effectively depict hidden information from figures in scientific publications. More details are presented in Sect. 5.1.

3.2 System overview

As illustrated in Fig. 1, VIStory workflow mainly consists of three stages: (1) *Data Extraction*, (2) *Data Management*, and (3) *Interactive Visual Exploration*. In *Data Extraction* stage, we extract figures and metadata from scientific publications, using an automatic figure extraction method (Sect. 4.1). This yields a total of 11,568 figures, and metadata including authors, venues, and keywords, from 1171 IEEE VIS publications 2009–2018. We also compute figure attributes including size, aspect ratio, and median color for each publication figure. Next, in the *Data Management* stage, we organize all extracted figures according to publication metadata and figure attributes in a nested table structure (Sect. 4.3). The structure enables quick query on publication metadata to support faceted exploration. Last, in *Interactive Visual Exploration* stage, we design *Faceted View*, *Storyboard View*, and *Endgame View* to facilitate exploration of the visual information. The interface is a web-based implementation. A demonstration of VIStory for the IEEE VIS publications can be found at: <https://dongoa.github.io/VIStory/>.

4 Modeling publication figure

This section first describes automatic methods for extracting figures (Sect. 4.1), followed by description of data characteristics (Sect. 4.2) and a nested table data structure for improving querying efficiency (Sect. 4.3).

4.1 Automatic figure extraction

Though we have crawled all the paper, it remains challenging to extract figures in an automatic way (*R1*). We develop an automatic figure extraction method with the following steps.

1. We first convert a PDF paper to JPG images using ghostscript,¹ and to an XML file using pdftohtml.² The XML file records ordered text boxes $\{B_i\}$ and their corresponding attributes of position (x , y), width (w), and height (h).

¹ www.ghostscript.com.

² pdftohtml.sourceforge.net.

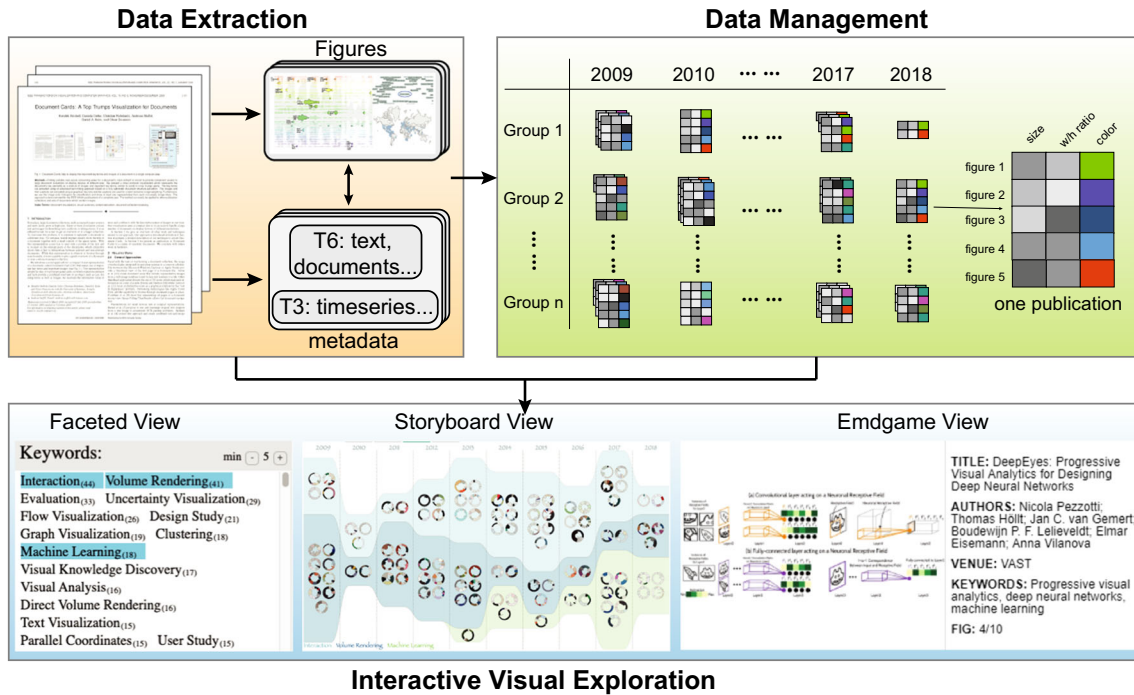


Fig. 1 VIStory workflow mainly consists of three stages: *Data Extraction*, *Data Management*, and *Interactive Visual Exploration*

2. We search for keywords of *Fig.* and *Figure* appearing as the first word of a text box, which indicate either figure captions or descriptions. Here, we make a reasonable heuristic that figure captions are placed below figures. Thus, if attribute y_i of a text box B_i to $(y_{i-1} + h_{i-1})$ of the previous box B_{i-1} is small, we regard the text box as text description.
3. After identifying a text box B_i as figure caption, we can determine a figure’s position in y-dimension as $(y_{i-1} + h_{i-1}, y_i)$. Position in x-dimension is determined by x_i and w_i . A two-column figure is identified if x_i is less than while $(x_i + w_i)$ is larger than half page width.
4. We use the identified positions to extract a figure in the corresponding JPG image. Lastly, we crop out background by identifying the minimum bounding box of pixels in different colors with the background color.

We measured accuracy of the automatic figure extraction method by comparing with figures extracted in a manual manner. We randomly selected 30 papers from the publication dataset, and manually cropped all the figures in each paper. First, we compared the accuracy regarding number of figures—the automatic method produces exactly the same number of figures with the manual cropping. Second, we compared the accuracy regarding cropping bounding box. Here, we only compared the aspect ratios between pairwise figures extracted by the two methods, since sizes of manually extracted figures are affected by screen size—the accuracy is 0.997.

Figure 2 presents average number of figures in the collected VAST, InfoVis, and SciVis publications. We can identify that most publications include 5–15 figures, while SciVis publications tend to have a slightly higher mean. Nevertheless, there are also several abnormalities. A VAST and InfoVis paper (highlighted by red circles) include over 30 figures, which is about three times than other papers.

4.2 Data characteristics

The IEEE VIS dataset (Isenberg et al. 2017a) also records multi-faceted metadata for each publication, including *venue* (i.e., VAST, InfoVis, SciVis), *publication year*, *authors*, and *keywords*. The metadata can reveal many interesting knowledge. For instance, we can figure out researchers who are active in the field by *author*, or find out what topics are becoming popular by *keyword*. Thus, we decide to support interactive

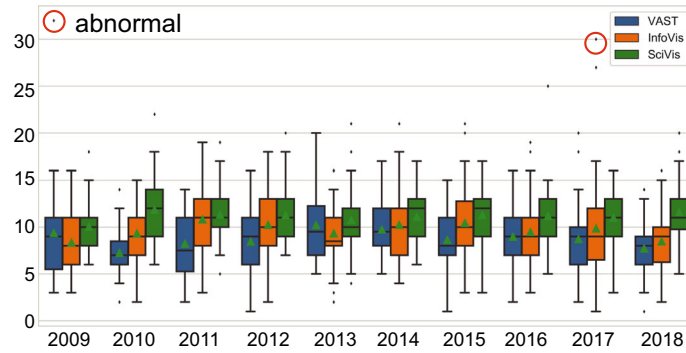


Fig. 2 Number of figures included in VAST, InfoVis, and SciVis proceedings in 2009–2018

exploration of the dataset using the metadata. We regard *publication year* as a key factor for depicting trends over time, thus it is fixed as a factor during interactive exploration. For the computer vision publications, we consider *num. of citations* as an important factor for reflecting the research impacts. Some publications exceeding 1,000 citations are of specific interest.

Besides publication metadata, we would like to further explore attributes of a figure F .

- *Figure size* (F_{size}). We measure figure size by multiplying figure width F_w and height F_h , i.e., $F_{size} = F_w \times F_h$.
- *Aspect ratio* (λ). We measure aspect ratio of a figure, i.e., $\lambda = F_w/F_h$.
- *Median color* (F_{color}). We identify median color of a figure defined as centroid of a cube representing all the colors enclosed by the cube (Heckbert 1982), which can be efficiently quantified using median-cut algorithm.

These attributes reveal intrinsic properties of visualizations as color and shape are pre-attentive visual stimuli (Rodrigues et al. 2006). In addition, it can also benefit paper writing, as researchers would like to know how much space and what w/h ratio is suitable. Notice that, image size and w/h ratio are scalar values, while color is represented as a vector of red, green, blue values.

4.3 Nested table structure

The processed data exhibit properties of *multi-faceted* (publication metadata) and *multi-dimensional* (figure attributes). Such complex data nature brings in challenges for accomplishing the requirement of $R2$. *multi-faceted exploration*. To overcome such challenge, we organize the data in a nested table structure, based on the universal relational model—‘one can place all data attributes into a table, which may then be decomposed into smaller tables as needed’ (Hawryszkiewicz 1984). Similar to that in (Shen et al. 2018), the table structure is illustrated in Fig. 1 top-right.

- We first group all publications by publication year in rows, and another attribute (venue, author, keyword, number of figures) in columns. By this, each cell consists of different numbers of publications.
- Each publication is further represented as a table recording figures as rows and figure attributes as columns. The number of rows is dynamical depending on number of figures, while column number is fixed to three for attributes of image size, aspect ratio, and color.

5 VISTory interface

Designing an intuitive visual interface is a key requirement of this work ($R3$). This section first summarizes carefully considered design rationales to fulfill the requirements, followed by detailed descriptions of components of VISTory.

5.1 Design rationales

We consider an intuitive visual design should meet the following rationales to fulfill the requirements:

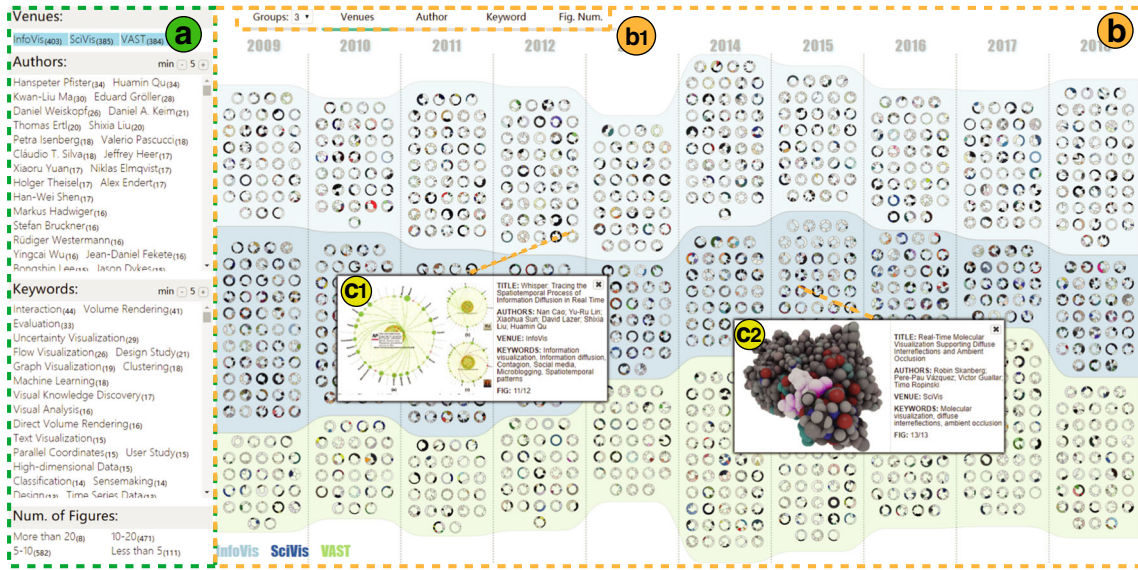


Fig. 3 VIStory interface for exploring a collection of scientific publications. **a** The Faceted View enables efficient query of publications through multi-faceted metadata of venues, authors, and keywords. **b** The queried publications are encoded as glyphs arranged in a themeriver layout to depict temporal trends in the Storyboard View. **c** The Endgame View presents a highlighted figure along with information of the publication

- *Complete*: The interface should support exploration of both publication metadata and figure attributes. The twofold perspective information complement each other in supporting high-level analytical tasks. For instance, to figure out what colors (*figure attributes*) are frequently used by a visualization expert (*publication metadata*).
- *Overview + Details*: No surprisingly, tremendous amount of figures will be collected from the publications. The system should provide overviews of the figures from different perspectives. Meanwhile, interactive techniques should be integrated to support details-on-demand exploration.
- *Faceted Browsing*: As described above, the publication metadata are faceted, i.e., composed of orthogonal sets of categories. To support *R2. Multi-faceted analysis*, the interface should allow users to manipulate the figures for analysis using faceted metadata, rather than projecting all figures into low-dimensions using MDS or t-SNE.

Based on these rationales, we finally come up with VIStory interface as shown in Fig. 3. The interface mainly consists of three view components: Faceted View (Fig. 3a), Storyboard View (Fig. 3b), and Endgame View (Fig. 3c).

5.2 Faceted view

Inspired by (Yee et al. 2003), we design Faceted View as shown in Fig. 3a to fulfill the rationale of *Faceted Browsing*. The view consists of configurable faceted panels corresponding to first level of metadata terms: For IEEE VIS publications, the panels are *Venues*, *Authors*, *Keywords*, and *Num. of Figures*, whilst for the second dataset, the panels are *Venues*, *Authors*, *Num. of Citations*, and *Num. of Figures*. Each panel is comprised of attributes of second level of metadata terms, e.g., *InfoVis*, *SciVis*, and *VAST* in the *Venues* panel for IEEE VIS publications, whilst *CVPR*, *ICCV*, *ECCV* for the second dataset. Specifically, we divide the numerical *Num. of Citations* into four ranges of *More than 1000*, *100–1000*, *50–100*, and *Less than 50*, and *Num. of Figures* into four ranges of *more than 20*, *10–20*, *5–10*, and *less than 5*. In this way, all faceted attributes are categorical. The attributes are sorted in descending order by the number of publications with the attribute. An exception is attributes in the fourth panel, which are sorted by number of figures. Notice that, there can be too many attributes in *Authors* and *Keywords* panels. We further add a minimum threshold controller to filter out attributes with fewer publications than the threshold.

To support *Overview + Details* rationale, the Faceted View enables visual query of publications for exploration by clicking on the user-interested attributes. Let denote the multi-faceted dataset of n publications as $\mathcal{P} = \{p_i\}_{i=1}^n$. We note the categorical facets as V (venues), A (authors), K (keywords), and N




Fig. 4 An example of query operations in VISTory

(num. of figures). For simplicity, we refer all of the facets as X unless stated explicitly. Let x_j as the j th attribute of facet X , $X(p_i)$ be the attribute value of facet X for publication p_i . Note that, $V(p_i)$ and $N(p_i)$ are single values, while $A(p_i)$ and $K(p_i)$ can be a vector of values as a publication may have multiple coauthors and keywords.

Let denote the list of publications with attribute x_j as \mathcal{P}_{x_j} . Thus, we have $\mathcal{P}_{x_j} = \{p \in \mathcal{P}, X(p) \in \{x_j\}\}$. In VISTory, users can select multiple attributes from the same or different facets. Figure 4 shows an example of visual query results made by two attributes V_1 and V_2 from facet V , and two attributes A_1 and A_2 from facet A .

- *Union*. In case, multiple attributes from the same facet are selected, query result is union of publications with the attributes, i.e., $\mathcal{P}_{x_{j_1}, x_{j_2}} = \{p \in \mathcal{P}, X(p) \in \{x_{j_1}, x_{j_2}\}\}$. For instance, when users select both A_1 and A_2 as illustrated in Fig. 4, the query result is $\{p_1, p_2, p_3, p_4\}$.
- *Intersection*. In case multiple attributes from different facets are selected, query result is intersection of publications with the attributes, i.e., $\mathcal{P}_{x_j, x'_j} = \{p \in \mathcal{P}, X(p) \in \{x_j\} \text{ and } X'(p) \in \{x'_j\}\}$. For instance, when users select both A_1 and V_1 , only publication p_1 will be queried.

By default, all attributes in *Venues* facet are selected (see Fig. 3a), i.e., all publications in the dataset are chosen for exploration.

5.3 Storyboard view

After a user selects certain attributes, a subset of publications $\mathcal{P}_s \subseteq \mathcal{P}$ is filtered for exploration. \mathcal{P}_s can be further grouped based on user-defined attribute of *Venues*, *Authors*, *Keyword*, or *Fig. Num.*, using the buttons in Fig. 3b1. Users can also control the number of groups to be visualized using the drop down selection list. After selecting the grouping attribute and number of groups, the relative information is presented in Storyboard View as shown in Fig. 3b. As a main view component in VISTory, the view employs the following intuitive visual designs.

5.3.1 Paper ring

To support *Complete* rationale, we need to depict multi-dimensional figure attributes, including *size*, *w/h ratio*, and *domain color*. We come up with a glyph of paper ring as shown in Fig. 5. Here, all figures in one publication are represented as arcs, which are arranged in a clockwise order corresponding to the figure order in the publication. For each figure, its *size* is encoded as arc length, *w/h ratio* encoded as arc height, and *domain color* as the arc color. Notice that, arc lengths indicate only relative sizes of figures in the same publication, but not absolute sizes across multiple publications. This work treats every publication equally, hence all paper rings share the same radius.

Figure 5 shows a paper ring glyph for a SciVis publication (Marino and Kaufman 2016). As the glyph depicts, there are in total 13 figures in the publication, and most figures exhibit brownish domain color. Obviously, the first figure occupies the largest size, while the third figure has the biggest w/h ratio. These two figures may reveal the main contributions of the publication. By examining the original figures, we can notice that the first figure includes two subfigures, while the third figure has three. These subfigures are arranged side-by-side, probably for comparative analysis.

5.3.2 Themeriver

There can be a numerous number of papers in \mathcal{P}_s , where a worst case can be 1171 when $\mathcal{P}_s = \mathcal{P}$. Remember that we treat year as a fixed analytical factor (Sect. 4.2). Thus, we employ *themeriver* (Havre et al. 2000)—a

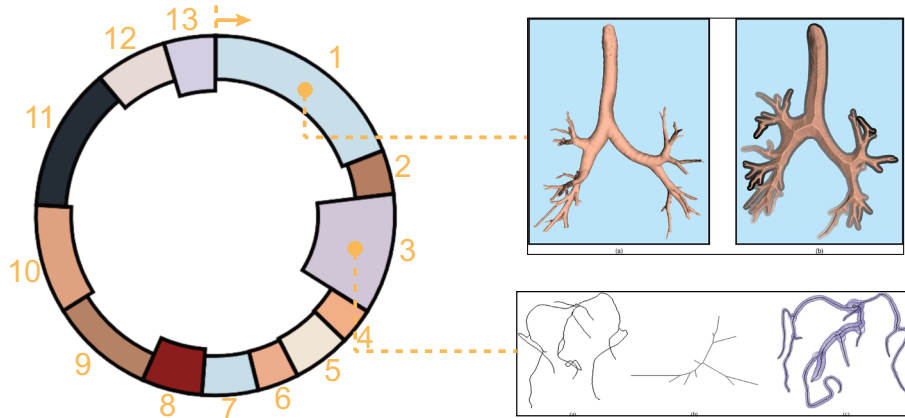


Fig. 5 Paper ring glyph for (Marino and Kaufman 2016): Arcs in clockwise layout depict figures in a paper, with arc length for figure size, arc height as w/h ratio, and color as domain color

classical visual representation for depicting temporal trends design, to arrange the paper rings. Here, the rendering canvas is first divided into 10 equal parts horizontally, corresponding to 10-years publication year. The river height corresponds to the number of publications. From Fig. 3, we can notice that the height of SciVis river (middle) is decreasing, while that of VAST river (bottom) is increasing.

5.3.3 Layout

Next, we need to position the paper rings in the themeriver in a meaningful way. Let denote a group of publications in one year as $\mathcal{P}_g := \{p_i\}_{i=1}^m$, where m indicates the number of publications in the group. We can extract a bounding box $\mathbf{B}_g := (cx_g, cy_g, w_g, h_g) \in R^2$ in the themeriver, where cx_g & cy_g indicate center position of \mathbf{B}_g , and w_g & h_g indicate its width and height, respectively. Our problem is to find $p_i := (cx_i, cy_i, r), \forall p_i \in \mathcal{P}_g$.

We develop a simple yet effective greedy algorithm to address this problem. The algorithm work as follows:

1. \mathbf{B}_g is first divided into 1 column and m rows, yielding $1 \times m$ grids and each grid can store one paper ring. We denote paper ring radius r as r_{g1} , and $r_{g1} = \min(w_g, \frac{h_g}{m})$.
2. We next divide \mathbf{B}_g into 2 columns and $\lceil m/2 \rceil$ rows, where $\lceil m/2 \rceil$ indicates the ceiling of $m/2$. In this way, we can derive $r_{g2} = \min(\frac{w_g}{2}, \frac{h_g}{\lceil m/2 \rceil})$.
3. We check condition $r_{g2} > r_{g1}$: if the condition is not met, we stop the process and return r_{g1} ; otherwise, we continue step 2 by increasing the column number until $r_{gn} < r_{g(n-1)}$ and return $r_{g(n-1)}$.
4. In the same way, we derive radii for all groups, and choose the minimum value as the final radius r .
5. After deciding r , we start from (cx_g, cy_g) and find a minimum bounding box that can pack all paper rings. In this way, paper rings in the same group are positioned close to each other and far from rings in other groups.

To better illustrate our solution, we give two examples shown in Fig. 6. The first example is storyboard for publications by *Hanspeter Pfister*, who got the most number of 34 IEEE VIS publications in 2009–2018 (together with Huamin Qu). The second example is storyboard for publications by top-3 authors of $\{Hanspeter Pfister, Huamin Qu, Kwan-Liu Ma\}$. Both views are grouped by *Venue* attribute. We select the first groups from both views, which contain one publication in (a) and four publications in (b). In (a), r_1 is chosen since $r_1 > r_2$; by contrast in (b), r_2 is chosen since $r_2 > r_1$ & $r_2 > r_3$.

5.4 Endgame view

To further support *Overview + Details* rationale, we further design Endgame View as shown in Fig. 3c1,c2. The view consists of two perspectives of information: First, the raw figure is presented on the left side; Second, the publication metadata, including *title*, *authors*, *venues*, *keywords*, and *order of the figure*, are presented on the right side. The view is connected to the center of its corresponding arc by a dashed line. It can be dragged around to avoid occlusion of important visuals. Multiple endgame views can be enabled at

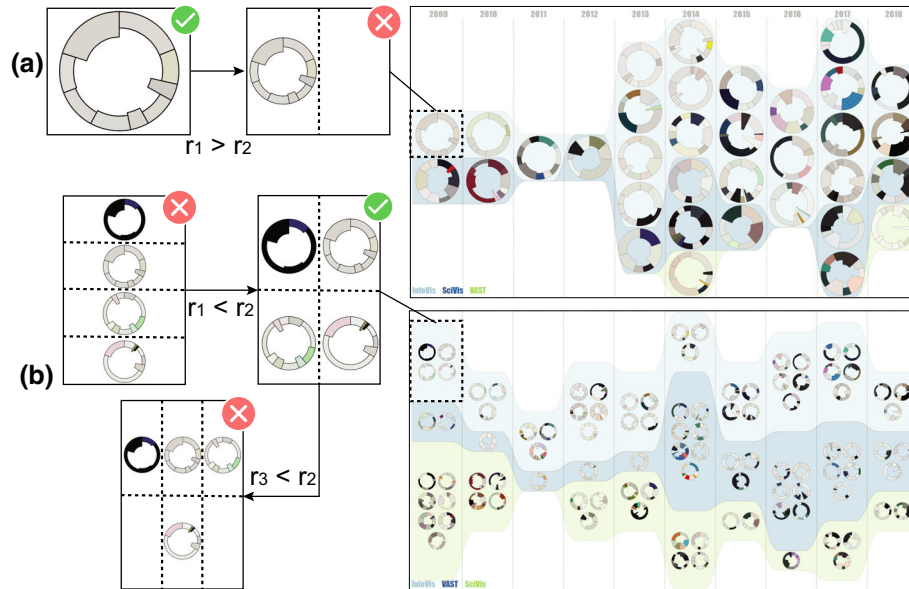


Fig. 6 A greedy algorithm for deciding radius of paper ring. r_1 is chosen in **a** while r_2 is chosen in **(b)**. See Sect. 5.3

the same time, such that to enable comparison. Taking Fig. 3c1,c2 for an example, it is obvious that (c1) presents a 2D abstract visualization with nodes and links, while (c2) is a scientific visualization with 3D visual cues.

6 Case study

We conduct case studies on two real-world datasets to demonstrate efficacy of VISTory. First, we experiment with representative visualizations from proceedings of renowned IEEE VIS conference (including VAST, InfoVis, and SciVis). To better understand trend of the field, we choose past 10-year papers published in 2009–2018. Thanks to a well-organized IEEE VIS dataset (Isenberg et al. 2017a), we can crawl all papers using provided digital object identifier (DOI). In total, we collect 1171 papers, of which 383 are VAST, 403 are InfoVis, and 385 are SciVis. Second, we extend the applicability to publications by a research group named Center for Multimedia Integrated Technologies at Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences. The team’s research focus on computer vision, multimedia, deep learning, etc. We select papers published by the team at top computer vision conferences of CVPR, ICCV, ECCV in 2007–2019, yielding a total of 43 papers, of which 21 are CVPR, 8 are ICCV and 14 are ECCV.

Below we present three usage scenarios: *author profile probe* (Sect. 6.1), *VIS trend analysis* (Sect. 6.2), and *research impact exploration* (Sect. 6.3).

6.1 Study 1: Author profile probe

VISTory can be applied to probe author profile, including number of publications and topics over years. To demonstrate this, we first filter publications made by active visualization researchers in mainland China. This is accomplished by selecting seven researchers of {Shixia Liu, Xiaoru Yuan, Yingcai Wu, Wei Chen, Weiwei Cui, Nan Cao, Yunhai Wang} (ordered by number of publications in past 10-year IEEE VIS conference) from *Author panel* in *Faceted View*. We then select author as grouping factor and set group number to 10, yielding a Storyboard View as in Fig. 7.

From the view, we can obtain several interesting discoveries.

- First, the view presents 10 author groups, meaning that there are three additional authors. We can identify they are *Huamin Qu*, *Hanqi Guo*, and *Mengchen Liu*, and they exhibit different patterns. (1) *Huamin Qu*: Surprisingly *Huamin Qu* obtains the most number of publications in the query result, indicating he had close collaborations with the selected seven researchers. (2) *Hanqi Guo*: We can notice

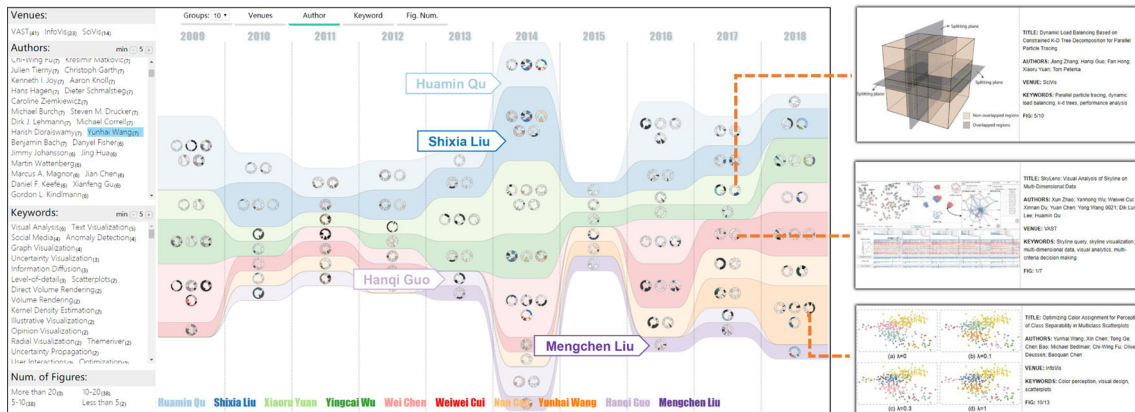


Fig. 7 Author profiles of top visualization researchers in mainland China who contributed most publications to IEEE VIS publications in 2009–2018, i.e., {*Shixia Liu, Xiaoru Yuan, Yingcai Wu, Wei Chen, Weiwei Cui, Nan Cao, Yunhai Wang*}. Notice that, *Huamin Qu, Hanqi Guo, and Mengchen Liu* are not in the selection list, but still appear in the view

- that the river of *Hanqi Guo* starts from 2010, and many paper rings are the same with those of *Xiaoru Yuan*. A quick examination reveals that they collaborated on seven publications in past 10 years. (3) *Mengchen Liu*: We can observe that *Mengchen Liu* stably contributed at least one publication every year starting from 2013, and all of them are in collaboration with *Shixia Liu*. The observations infer that *Hanqi Guo-Xiaoru Yuan* and *Mengchen Liu-Shixia Liu* were probably in supervisee–supervisor relationship.
- We can also observe peak and bottom publications years from the changes over time. In 2014, both *Shixia Liu* and *Xiaoru Yuan* contributed five publications, followed by four from *Wei Chen* and three from *Yingcai Wu*. In contrast, much fewer publications are made in 2011 and 2015.
 - Lastly, we would like to retrieve what topics the authors worked on, by clicking on figure arcs and examining the Endgame View. Here, we select three representative works as the insets by *Xiaoru Yuan* (top), *Weiwei Cui* (middle), and *Yunhai Wang* (bottom), which are published in *SciVis*, *VAST*, and *InfoVis*, respectively. The figures reflect different visualization techniques employed by the three publications. Scientific visualization (top inset) employs 3D visual representation to represent spatial attributes, visual analytics (middle inset) integrates coordinated multiple views to depict data from multiple perspectives, and information visualization (bottom inset) focuses on improving human’s visual perception on abstract 2D data.

6.2 Study 2: VIS trend analysis

VIStory can also be utilized to analyze trend of visualization topics, by exploring keyword attribute. Here, we filter publications by selecting three top keywords of {*Interaction, Volume Rendering, and Machine Learning*} from the *Keyword* panel in *Faceted View*. There are in total 44 publications with the keyword of *interaction*, 41 for *volume rendering*, and 18 for *machine learning*. Keyword is also selected as grouping factor and group number is set to 3, yielding a *Storyboard View* presented in Fig. 8 (left).

Through the storyboard, we can observe different trend patterns of the topics in past 10-year IEEE VIS conference. First, number of publications on *interaction* remains relatively stable, with several publications accepted each year. *Interaction* is a fundamental component for interactive visualization, thus many studies were conducted to improve interaction efficiency. Second, *volume rendering* becomes less popular, as the number of publications is decreasing. This is probably because volume rendering has been exhaustively studied since the beginning of scientific visualization. In contrast, we can observe that more publications on *machine learning* were accepted in the last two years. A heuristic is that many visual analytics systems have been developed to open the black box of deep learning techniques.

To verify the heuristic, we dis-select *interaction* and *volume rendering*, and change grouping factor to venue. This results in a *Storyboard View* as presented in Fig. 8 (middle). A first glimpse shows that most machine learning publications are in *VAST* conference, while only two in *InfoVis* and one in *SciVis*. Deeper examination by viewing an *Endgame View* (top inset) reveals that the work is to visualize training process




Fig. 8 Exploring trends of visualization topics of *Interaction*, *Volume Rendering*, and *Machine Learning* in 2009–2018. Left: Number of publications on interaction remains stable, on volume rendering is decreasing, while on machine learning gets popular. Center: Closer examination of publications on machine learning shows that most publications are in VAST conference. Right: Top Endgame View shows that the VAST paper is to explain CNN, while the bottom view shows a InfoVis paper utilizing machine learning to facilitate graph layout

of a convolutional neural network (CNN). In contrast, another Endgame View (bottom inset) shows that the work utilized machine learning to facilitate graph layout. The findings indicate that the heuristic is reasonable.

6.3 Study 3: Research impact exploration

We conduct another study on the second dataset, *i.e.*, research papers published at CVPR, ICCV, and ECCV by a multimedia team. The team was established in 2007, so we select the period from 2007–2019. Figure 9 (left) presents an overview of the publications. From the overview, we can notice that: First, the team achieved much better performance at the three major computer vision conferences since 2013. This indicates that the team grows up since then, and the team has contributed much to the trend of applying deep learning in vision tasks. Second, the team publish more in CVPR than the other two conferences. There are no overlap between ICCV and ECCV publications in each year. This is probably because ICCV is held every other year. Third, we can see that most arcs exhibit either white or dark gray colors. In comparison, visualization figures (see Figs. 7, 8) present more diverse colors. This reflects the differences in color usage by the two fields.

Next, we would like to explore high-impact publications by the team. This is done by filtering publications with over 1000 citations. Fig. 9(center) shows the results. We can see that there are a total of five publications, out of which four are published in ECCV and one in CVPR. All papers were published before 2017, as it takes time for paper citing. Actually, there are several other papers after 2017 with hundreds of citations. By clicking on the figure arcs, we can exploit the details by checking the Endgame views (Fig. 9(right)). The top view shows the corresponding paper is about action recognition from videos, while the bottom view shows the paper is on image super-resolution. Both papers are pioneering works on the topics.




Fig. 9 Exploring research impacts of computer vision papers by a multimedia group. Left: overview of all publications by the team at CVPR, ICCV, and ECCV conferences. Center: Exploration of high-impact publications by filtering publications with more than 1000 citations. Right: Endgame Views show details of these high-impact publications

7 User study

To evaluate the effectiveness of *VIStory* in assisting users in exploring scientific publications, we conducted a user study that compares *VIStory* with *SurVis* (Beck et al. 2016)—a state-of-the-art visual literature survey system. Similar to *VIStory*, *SurVis* also enables the examination of scientific publications from multiple facets, e.g., authors and keywords. Yet, *SurVis* lacks an overview of the visual information. For fair comparison, we replicated a *SurVis* implementation of the IEEE VIS publications using source code (<https://github.com/fabian-beck/survis>) provided by the authors. The reimplementation can be found on <https://xiii.bitbucket.io/>.

Participants We recruited 20 participants (6 females, 14 males) between the ages of 23–25 (age: 24.25 ± 0.91). All participants are graduate students with backgrounds in computer science. They are familiar with digital libraries, e.g., *Google Scholar* and *Microsoft Academic*, in retrieving literature.

Experiment setting and procedure We prepared 10 multiple-choice questions regarding information in scientific publication, e.g., “Which keyword (*Interaction*, *Volume Rendering*, *Machine Learning*) is an emerging topic over 2009–2018?” and “Which authors lead in ‘*Uncertainty Visualization*’ publications during 2009–2018? ”. The questions were well chosen to cover both metadata and visual information of the scientific publications; see Supplementary Table S1. An optional answer ‘Cannot find the answer’ is provided for each question, in case, the participant felt the answer is not available. Each participant was asked to find answers using both *VIStory* and *SurVis*. To minimize learning effects, the order of systems was randomly assigned to each participant. Due to the COVID-19 pandemic, the study was performed virtually.

Each participant first performed a pre-study background questionnaire, followed by an introduction of the functionalities of both systems. We next allowed the participants to freely explore the systems until they felt comfortable using them. After that, the participants went through the questions using the systems. The participants were reminded to finish the questions as fast as they can, as completion time is an evaluation metrics. After finishing the questions, the participants were asked to complete a questionnaire.

Hypothesis: *VIStory* provides an overview of the scientific publications by arranging paper glyphs in a themeriver layout, which is not available in *SurVis*. Hence, we expect *VIStory* would be more efficient, i.e., taking less completion time, than *SurVis*. **Result:** We collected in total 40 (20 participants \times 2 systems) experiment results. Before the analysis, we first confirmed that all results of completion time follow a normal distribution using a Shapiro–Wilk test. We then performed a one-way ANOVA on two groups of experiment results for the systems. Completion time for the questions with *VIStory* is on average 142.75s less than that with *SurVis* ($p < 0.05$) (see Fig. 10 (left)). The results confirmed the hypothesis. Notice that, the average number of correct answers for *VIStory* (9.2) is higher than that for *SurVis* (6.0) (see Fig. 10 (left)). This is because *SurVis* shows no visual information that is required for some questions. We observed that the participants quickly skipped the questions by choosing ‘Cannot find the answer’ option. Hence, *VIStory* would be even more efficient than *SurVis* in terms of completion time per correct answer.

Feedback: Qualitative feedbacks using 7-point Likert scale questions were collected from the participants after the experiments. Figure 10 (right) presents a summary of the feedbacks. More details are presented in Supplementary Table S2. For the interface design, the participants had a positive impression. They all agreed that (1) the *paper ring* is intuitive (Q1) (mean = 5.55, SD = 1.02); (2) the *themeriver* layout depicts temporal trends well (Q2) (mean = 6.15, SD = 0.79); and (3) the faceted exploration facilitate multi-

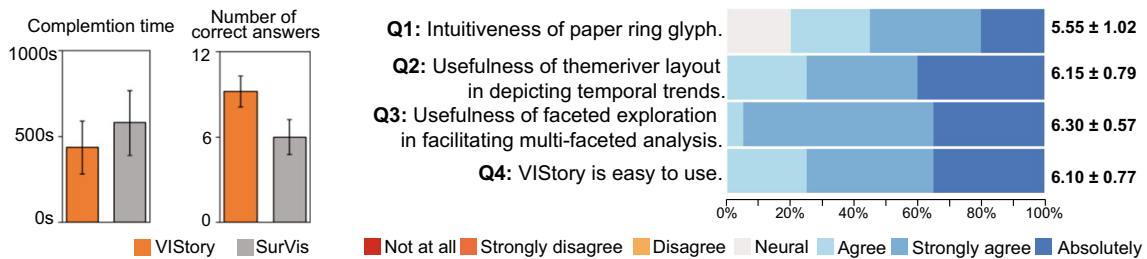


Fig. 10 User study results. Left: Quantitative results of means and standard errors of completion time and number of correct answers using *VIStory* and *Survis*. Right: Qualitative feedbacks on usability of *VIStory*. The rightmost column denotes Mean \pm SD

faceted analysis (Q3) (mean = 6.30, SD = 0.57). For the system usability, all participants consider *VIStory* is easy to use (Q4) (mean = 6.10, SD = 0.77).

In the free-form question on future improvements (Q5), the participants gave several fruitful suggestions. First, some participants noticed that many paper rings are in light colors, making it difficult to distinguish with the background themeriver colors. The participants tended to examine endgame views of those paper rings in dark colors. We explained that the ring color is corresponding to the median color of the figure, and the participants suggested to encode some other metrics, such as color histograms and variance. Second, the participants also suggested to integrate the functionality of query by image in the interface. Several participants studying computer vision noted that figure plagiarism detection is getting more attention recently, and to find similar visualization designs would be an interesting topic. This will make *VIStory* more useful, and we plan to realize it in the near future.

8 Limitations and future work

The case studies demonstrate the efficacy of *VIStory* in probing author profiles, understanding visualization trends, and exploring research impacts. These information can benefit real-world applications, e.g., to help students find suitable supervisors, and to help researchers find hot topics. Nevertheless, there are still some limitations of our system.

First, this analyses are conducted on past 10-year IEEE VIS publications and publications by a specific research team. The information only covers a small amount of works in visualization and computer vision. For instance, volume rendering as a pioneering visualization topic has now been widely used for visualizing medical images and flow simulations. Many studies on volume rendering have been published on other venues such as *IEEE Transactions on Medical Imaging (IEEE TMI)*. Similarly, IEEE VIS publications only count a small portion of outcomes of an author. Many publications, such as those in *IEEE Transactions on Visualization and Computer Graphics (IEEE TVCG)*, are not counted here. In this sense, we can only claim that study 1 reveals author profiles, and study 2 indicates trends of visualization topics in IEEE VIS conference. Nevertheless, we regard this as a common limitation for similar studies using only IEEE VIS publications (e.g., Isenberg et al. 2017a, b).

A feasible solution to address the limitation is to incorporate more data for analysis, e.g. other publications in IEEE TVCG, EuroVis, PacificVis, and VINCI. In this way, a more complete overview of the visualization field can be achieved. However, this can cause another limitation regarding scalability of the system. Experiments reveal that paper rings in Storyboard View becomes too small to be observable when the total number of publications reaches 1000 (see Fig. 3 for an example). Though the scalability issue can be mitigated through filtering interactions, we would like to examine more visual design alternatives. A feasible solution here is to employ advanced semantic image projection methods, which has been shown effective for handling millions of images (Xie et al. 2018). To integrate semantic image projection with faceted visual interaction (Yee et al. 2003) would be an interesting direction.

Besides, there are several promising directions for our future work. First, we spent much time on extracting figures from visualization publications. We would like to make it open for future researches, e.g., to extract visualization-related image metrics using deep learning techniques. We also call for collaborations on enriching the dataset, such as to manually label all the figures. We will soon make *VIStory* system open to the public. Second, with the automatic figure extraction method, we can easily harvest more figures from scientific publications. We plan to do so on publications in past 10-year IEEE TVCG, EuroVis, and PacificVis. Lastly, we would like to continue working on the visual interface to incorporate more analytical features and improve the system scalability.

9 Conclusion

This paper presents *VIStory*, a storyboard that supports interactive exploration of visual information collected from scientific publications. The benefit of this work is prominent: First, we suggest visual information as a new perspective for exploring scientific publications. To illustrate the usability of this approach, we curate a new dataset of 11,568 figures from ten-year IEEE VIS publications in 2009–2018 by an automatic figure extraction method. We plan to release the dataset to facilitate future research. Second, we develop *VIStory*—a literature survey system that assists users in exploring the visual information. The

system integrates a nested table structure to support multi-faceted analysis, and design of intuitive paper rings arranged in a themeriver layout to promote intuitive visual exploration. Third, we conduct three case studies and a user study that demonstrate the effectiveness of VISTORY in helping users address practical needs such as to probe author profile, to identify temporal trends, and to explore research impact.

Acknowledgements The authors wish to thank the experiment participants for their valuable feedbacks, the authors of SurVis for their open-source code, and the reviewers for their constructive comments. This work was supported by National Natural Science Foundation of China (Nos. 61972388, 61802388).

References

- Beck F, Koch S, Weiskopf D (2016) Visual analysis and dissemination of scientific literature collections with SurVis. *IEEE Trans Vis Comput Graph* 22(1):180–189. <https://doi.org/10.1109/TVCG.2015.2467757>
- Borkin MA, Vo AA, Bylinskii Z, Isola P, Sunkavalli S, Oliva A, Pfister H (2013) What makes a visualization memorable? *IEEE Trans Vis Comput Graph* 19(12):2306–2315. <https://doi.org/10.1109/TVCG.2013.234>
- Borkin MA, Bylinskii Z, Kim NW, Bainbridge CM, Yeh CS, Borkin D, Pfister H, Oliva A (2016) Beyond memorability: visualization recognition and recall. *IEEE Trans Vis Comput Graph* 22(1):519–528. <https://doi.org/10.1109/TVCG.2015.2467732>
- Chen M, Ebert D, Hagen H, Laramée RS, van Liere R, Ma KL, Ribarsky W, Scheuermann G, Silver D (2009) Data, information, and knowledge in visualization. *IEEE Comput Graph Appl* 29(1):12–19. <https://doi.org/10.1109/MCG.2009.6>
- Chen X, Zeng W, Lin Y, Al-manee HM, Roberts J, Chang R (2020) Composition and configuration patterns in multiple-view visualizations. *IEEE Trans Vis Comput Graph*. Accepted
- Chou JK, Yang CK (2011) PaperVis: literature review made easy. *Comput Graphics Forum* 30(3):721–730. <https://doi.org/10.1111/j.1467-8659.2011.01921.x>
- Dou W, Yu L, Wang X, Ma Z, Ribarsky W (2013) Hierarchical topics: visually exploring large text collections using topic hierarchies. *IEEE Trans Vis Comput Graph* 19(12):2002–2011. <https://doi.org/10.1109/TVCG.2013.162>
- Eler DM, Nakazaki MY, Paulovich FV, Santos DP, Andery GF, Oliveira MCF, Batista Neto J, Minghim R (2009) Visual analysis of image collections. *Vis Comput* 25(10):923–937. <https://doi.org/10.1007/s00371-009-0368-7>
- Federico P, Heimerl F, Koch S, Miksch S (2017) A survey on visual approaches for analyzing scientific literature and patents. *IEEE Trans Vis Comput Graph* 23(9):2179–2198. <https://doi.org/10.1109/TVCG.2016.2610422>
- Görg C, Liu Z, Kihm J, Choo J, Park H, Stasko J (2013) Combining computational analyses and interactive visualization for document exploration and sensemaking in Jigsaw. *IEEE Trans Vis Comput Graph* 19(10):1646–1663. <https://doi.org/10.1109/TVCG.2012.324>
- Havre S, Hetzler B, Nowell L (2000) ThemeRiver: visualizing theme changes over time. In: *Proceedings of IEEE symposium on information visualization*, pp 115–123. <https://doi.org/10.1109/INFVIS.2000.885098>
- Hawryszkiewicz IT (1984) *Database analysis and design*. Science Research Associates
- Heckbert P (1982) Color image quantization for frame buffer display. *SIGGRAPH Comput Graph* 16(3):297–307. <https://doi.org/10.1145/965145.801294>
- Heimerl F, Han Q, Koch S, Ertl T (2016) CiteRivers: Visual analytics of citation patterns. *IEEE Trans Vis Comput Graph* 22(1):190–199. <https://doi.org/10.1109/TVCG.2015.2467621>
- Isenberg P, Heimerl F, Koch S, Isenberg T, Xu P, Stolper CD, Sedlmair M, Chen J, Möller T, Stasko J (2017a) Vispubdata.org: a metadata collection about IEEE visualization (VIS) publications. *IEEE Trans Vis Comput Graph* 23(9):2199–2206. <https://doi.org/10.1109/TVCG.2016.2615308>
- Isenberg P, Isenberg T, Sedlmair M, Chen J, Möller T (2017b) Visualization as seen through its research paper keywords. *IEEE Trans Vis Comput Graph* 23(1):771–780. <https://doi.org/10.1109/TVCG.2016.2598827>
- Jänicke H, Chen M (2010) A saliency-based quality metric for visualization. *Comput Graph Forum* 29(3):1183–1192. <https://doi.org/10.1111/j.1467-8659.2009.01667.x>
- Joia P, Coimbra D, Cuminato JA, Paulovich FV, Nonato LG (2011) Local affine multidimensional projection. *IEEE Trans Vis Comput Graph* 17(12):2563–2571. <https://doi.org/10.1109/TVCG.2011.220>
- Keim D, Andrienko G, Fekete JD, Görg C, Kohlhammer J, Melançon G (2008) *Visual analytics: definition, process, and challenges*, vol 4950. *Lecture notes in computer science*. Springer, Berlin, pp 154–175
- Kucher K, Kerren A (2015) Text visualization techniques: taxonomy, visual survey, and community insights. In: *2015 IEEE Pacific visualization symposium (PacificVis)*, pp 117–121. <https://doi.org/10.1109/PACIFICVIS.2015.7156366>
- Latif S, Beck F (2018) VIS author profiles: interactive descriptions of publication records combining text and visualization. *IEEE Trans Vis Comput Graph*, pp 1–1. <https://doi.org/10.1109/TVCG.2018.2865022>
- Liang Y, Wang X, Zhang S, Hu S, Liu S (2018) PhotoRecomposer: interactive photo recombination by cropping. *IEEE Trans Vis Comput Graph* 24(10):2728–2742. <https://doi.org/10.1109/TVCG.2017.2764895>
- Marino J, Kaufman A (2016) Planar visualization of treelike structures. *IEEE Trans Vis Comput Graph* 22(1):906–915. <https://doi.org/10.1109/TVCG.2015.2467413>
- Matzen LE, Haass MJ, Divis KM, Wang Z, Wilson AT (2018) Data visualization saliency model: a tool for evaluating abstract data visualizations. *IEEE Trans Vis Comput Graph* 24(1):563–573. <https://doi.org/10.1109/TVCG.2017.2743939>
- Plant W, Schaefer G (2011) *Visualisation and browsing of image databases*. Springer, Berlin, pp 3–57

- Rodrigues JF, Traina AJM, Oliveira MCFd, Traina C (2006) Reviewing data visualization: an analytical taxonomical study. In: Proceedings of tenth international conference on information visualisation, pp 713–720. <https://doi.org/10.1109/IV.2006.94>
- Rooij O, Jv Wijk, Worring M (2010) MediaTable: interactive categorization of multimedia collections. *IEEE Comput Graph Appl* 30(5):42–51. <https://doi.org/10.1109/MCG.2010.66>
- Schulz HJ (2011) Treevis.net: a tree visualization reference. *IEEE Comput Graph Appl* 31(6):11–15. <https://doi.org/10.1109/MCG.2011.103>
- Shen Q, Zeng W, Ye Y, Mueller Arisona S, Schubiger S, Burkhard R, Qu H (2018) Streetvizor: visual exploration of human-scale urban forms based on street views. *IEEE Trans Vis Comput Graph* 24(1):1004–1013. <https://doi.org/10.1109/TVCG.2017.2744159>
- Shneiderman B (1996) The eyes have it: a task by data type taxonomy for information visualizations. In: Proceedings of IEEE symposium on visual languages, IEEE, pp 336 – 343. <https://doi.org/10.1109/VL.1996.545307>
- Strobel H, Oelke D, Rohrdantz C, Stoffel A, Keim DA, Deussen O (2009) Document cards: a top trumps visualization for documents. *IEEE Trans Vis Comput Graph* 15(6):1145–1152. <https://doi.org/10.1109/TVCG.2009.139>
- Tan L, Song Y, Liu S, Xie L (2012) ImageHive: interactive content-aware image summarization. *IEEE Comput Graph Appl* 32(1):46–55. <https://doi.org/10.1109/MCG.2011.89>
- Tominski C, Aigner W (2015) The TimeViz browser. <http://survey.timeviz.net/>. Accessed 18 May 2020
- Tory M, Möller T (2004) Rethinking visualization: a high-level taxonomy. In: IEEE symposium on information visualization, pp 151–158. <https://doi.org/10.1109/INFVIS.2004.59>
- Van Der Corput P, & van Wijk JJ (2016) ICLIC: interactive categorization of large image collections. In: Proceedings of IEEE Pacific visualization symposium (PacificVis), pp 152–159. <https://doi.org/10.1109/PACIFICVIS.2016.7465263>
- van der Corput P, van Wijk JJ (2017) Comparing personal image collections with PICTuReVis. *Comput Graph Forum* 36(3):295–304. <https://doi.org/10.1111/cgf.13188>
- Wu Y, Pitipornvivat N, Zhao J, Yang S, Huang G, Qu H (2015) egoSlider: visual analysis of egocentric network evolution. *IEEE Trans Vis Comput Graph* 22(1):260–269. <https://doi.org/10.1109/TVCG.2015.2468151>
- Xie X, Cai X, Zhou J, Cao N, Wu Y (2018) A semantic-based method for visualizing large image collections. *IEEE Trans Vis Comput Graph* pp 1-1. <https://doi.org/10.1109/TVCG.2018.2835485>
- Yang J, Fan J, Hubball D, Gao Y, Luo H, Ribarsky W, Ward M (2006) Semantic image browser: bridging information visualization with automated intelligent image analysis. In: Proceedings of IEEE symposium on visual analytics science and technology, pp 191–198. <https://doi.org/10.1109/VAST.2006.261425>
- Yee KP, Swearingen K, Li K, Hearst M (2003) Faceted metadata for image search and browsing. In: Proceedings of the SIGCHI conference on human factors in computing systems, ACM 642681, pp 401–408. <https://doi.org/10.1145/642611.642681>
- Zeng W, Shen Q, Jiang Y, Telea A (2019) Route-aware edge bundling for visualizing origin-destination trails in urban traffic. *Comput Graph Forum* 38(3):581–593. <https://doi.org/10.1111/cgf.13712>