

Big Data Visualization Tools ^{*}

Nikos Bikakis
ATHENA Research Center, Greece

1 Synonyms

Visual exploration; Interactive visualization; Information visualization; Visual analytics; Exploratory data analysis.

2 Definition

Data visualization is the presentation of data in a pictorial or graphical format, and a *data visualization tool* is the software that generates this presentation. Data visualization provides users with intuitive means to interactively explore and analyze data, enabling them to effectively identify interesting patterns, infer correlations and causalities, and supports sense-making activities.

3 Overview

Exploring, visualizing and analysing data is a core task for data scientists and analysts in numerous applications. *Data visualization*² [1] provides intuitive ways for the users to interactively explore and analyze data, enabling them to effectively identify interesting patterns, infer correlations and causalities, and support sense-making activities.

^{*} This article appears in *Encyclopedia of Big Data Technologies*, Springer, 2018

² Throughout the article, terms *visualization* and *visual exploration*, as well as terms *tool* and *system* are used interchangeably.

The Big Data era has realized the availability of a great amount of *massive* datasets that are *dynamic*, *noisy* and *heterogeneous* in nature. The level of difficulty in transforming a data-curious user into someone who can access and analyze that data is even more burdensome now for a great number of users with little or no support and expertise on the data processing part. Data visualization has become a major research challenge involving several issues related to data storage, querying, indexing, visual presentation, interaction, personalization [2, 3, 4, 5, 6, 7, 8, 9].

Given the above, modern visualization and exploration systems should effectively and efficiently handle the following aspects.

- *Real-time Interaction*. Efficient and scalable techniques should support the interaction with billion objects datasets, while maintaining the system response in the range of a few milliseconds.
- *On-the-fly Processing*. Support of on-the-fly visualizations over large and dynamic sets of volatile raw (i.e., not preprocessed) data is required.
- *Visual Scalability*. Provision of effective data abstraction mechanisms is necessary for addressing problems related to visual information overloading (a.k.a. overplotting).
- *User Assistance and Personalization*. Encouraging user comprehension and offering customization capabilities to different user-defined exploration scenarios and preferences according to the analysis needs are important features.

4 Visualization in Big Data Era

This section discusses the basic concepts related to Big Data visualization. First, the limitations of traditional visualization systems are outlined. Then, the basic characteristics of data visualization in the context of Big Data era are presented. Finally, the major prerequisites and challenges that should be addressed by modern exploration and visualization systems are discussed.

4.1 Traditional Systems

Most *traditional exploration and visualization systems* cannot handle the size of many contemporary datasets. They restrict themselves to dealing with *small* dataset sizes, which can be easily handled and analysed with conventional data management and visual explorations techniques. Further, they operate in an *offline* way, limited to accessing *preprocessed* sets of *static* data.

4.2 Current Setting

On the other hand, nowadays, the *Big Data era* has made available large numbers of *very big* datasets, that are often *dynamic* and characterized by high *variety* and *volatility*. For example, in several cases (e.g., scientific databases), new data constantly arrive (e.g., on a daily/hourly basis); in other cases, data sources offer query or API endpoints for online access and updating. Further, nowadays, an increasingly large number of *diverse users* (i.e., users with different preferences or skills) explore and analyze data in a plethora of *different scenarios*.

4.3 Modern Systems

Modern systems should be able to efficiently handle *big dynamic* datasets, operating on machines with limited computational and memory resources (e.g., laptops). The dynamic nature of nowadays data (e.g., stream data), hinders the application of a preprocessing phase, such as traditional database loading and indexing. Hence, systems should provide *on-the-fly processing* over large sets of raw data.

Further, in conjunction with performance issues, modern systems have to address challenges related to *visual presentation*. Visualizing a large number of data objects is a challenging task; modern systems have to “*squeeze a billion records into a million pixels*” [3]. Even in small datasets, offering a dataset *overview* may be extremely difficult; in both cases, *information overloading* (a.k.a. overplotting) is a common issue. Consequently, a basic requirement of modern systems is to effectively support *data abstraction* over enormous numbers of data objects.

Apart from the aforementioned requirements, modern systems must also satisfy the diversity of *preferences* and *requirements* posed by different *users* and *tasks*. Modern systems should provide the user with the ability to customize the exploration experience based on her preferences and the individual requirements of each examined task. Additionally, systems should automatically adjust their parameters by taking into account the *environment setting* and *available resources*; e.g., screen resolution/size, available memory.

5 Systems and Techniques

This section presents how state-of-the-art approaches from Data Management and Mining, Information Visualization and Human-Computer Interaction communities attempt to handle the challenges that arise in the Big Data era.

5.1 Data Reduction

In order to handle and visualize large datasets, modern systems have to deal with information overloading issues. Offering *visual scalability* are crucial in Big Data visualization. Systems should provide efficient and effective abstraction and summarisation mechanisms. In this direction, a large number of systems use *approximation techniques* (a.k.a. *data reduction techniques*), in which abstract sets of data are computed. Considering the existing approaches, most of them are based on: (1) *sampling* and *filtering* [10, 11, 12, 13, 14] and/or (2) *aggregation* (e.g., binning, clustering) [15, 16, 17, 18, 19, 20].

5.2 Hierarchical Exploration

Approximation techniques are often defined in a hierarchical manner [15, 16, 19, 20], allowing users to explore data in different levels of detail (e.g., hierarchical aggregation).

*Hierarchical approaches*³ allow the visual exploration of very large datasets in multiple levels, offering both an overview, as well as an intuitive and effective way for finding specific parts within a dataset. Particularly, in hierarchical approaches, the user first obtains an *overview* of the dataset before proceeding to data exploration operations (e.g., roll-up, drill-down, zoom, filtering) and finally retrieving *details* about the data. Therefore, hierarchical approaches directly support the visual information seeking mantra “*overview first, zoom and filter, then details on demand*” [21]. Hierarchical approaches can also effectively address the problem of information overloading as they adopt approximation techniques.

Hierarchical techniques have been extensively used in large *graphs visualization*, where the graph is recursively decomposed into smaller sub-graphs that form a hierarchy of abstraction layers. In most cases, the hierarchy is constructed by exploiting *clustering* and *partitioning* methods [22, 23, 24, 25]. In other works, the hierarchy is defined with *hub-based* [26] and *density-based* [27] techniques. [28] supports *ad-hoc hierarchies* which are manually defined by the users. Different approaches have been adopted in [29, 30], where sampling techniques have been exploited. Other works adopt edge bundling techniques which aggregate graph edges to bundles [31, 32, 33, 34, 35, 36].

³ sometimes also referred as *multilevel*

5.3 *Progressive Results*

Data exploration requires real-time system's response. However, computing *complete results* over large (unprocessed) datasets may be extremely costly and in several cases *unnecessary*. Modern systems should *progressively* return *partial* and preferably *representative* results, as soon as possible.

Progressiveness can significantly improve efficiency in *exploration scenarios*, where it is common that users attempt to find something interesting without knowing what exactly they are searching for beforehand. In this case, users perform a sequence of operations (e.g., queries), where the result of each operation determines the formulation of the next operation. In systems where progressiveness is supported, in each operation, after inspecting the already produced results, the user is able to interrupt the execution and define the next operation, without waiting the exact result to be computed.

In this context, several systems adopt *progressive techniques*. In these techniques the results/visual elements are computed/constructed incrementally based on user interaction or as time progresses [16, 37, 38]. Further, numerous recent systems integrate incremental and approximate techniques. In these cases, approximate results are computed incrementally over progressively larger samples of the data [10, 12, 13].

5.4 *Incremental and Adaptive Processing*

The dynamic setting established nowadays hinders (efficient) data preprocessing in modern systems. Additionally, it is common in exploration scenarios that only a small fragment of the input data to be accessed by the user.

In situ data exploration [16, 39, 40, 41, 42, 43] is a recent trend, which aims at enabling on-the-fly exploration over large and dynamic sets of data, without (pre)processing (e.g., loading, indexing) the whole dataset. In these systems, *incremental* and *adaptive* processing and indexing techniques are used, in which small parts of data are processed incrementally “following” users' interactions.

5.5 *Caching and Prefetching*

Recall that, in exploration scenarios, a sequence of operations is performed and, in most cases, each operation is driven by the previous one. In this setting, *caching* and/or *prefetching* the sets of data that are likely to be accessed by the user in the near future can significantly reduce the response time [16, 38, 44, 45, 46, 47, 48]. Most of these approaches use prediction tech-

niques which exploit several factors (e.g., user behavior, user profile, use case) in order to determine the upcoming user interactions.

5.6 User Assistance

The large amount of available information makes it difficult for users to manually explore and analyze data. Modern systems should provide mechanisms that assist the user and reduce the effort needed on their part, considering the diversity of preferences and requirements posed by different users and tasks.

Recently, several approaches have been developed in the context of *visualization recommendation* [49]. These approaches recommend the most suitable visualizations in order to assist users throughout the analysis process. Usually, the recommendations take into account several factors, such as data characteristics, examined task, user preferences and behavior, etc.

Especially considering data characteristics, there are several systems that recommend the most suitable visualization technique (and parameters) based on the type, attributes, distribution, or cardinality of the input data [16, 50, 51, 52, 53, 54]. In a similar context, some systems assist users by recommending certain visualizations that reveal surprising and/or interesting data [55, 56, 57]. Other approaches provide visualization recommendations based on user behavior and preferences [58, 59]. Finally, systems provide recommendations and explanations regarding data trends and anomalies [60, 61].

6 Examples of Applications

Visualization techniques are of great importance in a wide range of application areas in the Big Data era. The volume, velocity, heterogeneity and complexity of available data make it extremely difficult for humans to explore and analyze data. Data visualization enables users to perform a series of analysis tasks that are not always possible with common data analysis techniques [64].

Major application domains for data visualization and analytics are *Physics* and *Astronomy*. Satellites and telescopes collect daily massive and dynamic streams of data. Using traditional analysis techniques, astronomers are able to identify noise, patterns and similarities. On the other hand, visual analytics can enable astronomers to identify unexpected phenomena and perform several complex operations, which are not feasible by traditional analysis approaches.

Another application domain is *atmospheric sciences* like *Meteorology* and *Climatology*. In this domain high volumes of data are collected from sensors

and satellites on a daily basis. Storing these data over the years results in massive amounts of data that have to be analyzed. Visual analytics can assist scientists to perform core tasks, such as climate factors correlation analysis, event prediction, etc. Further, in this domain, visualization systems are used in several scenarios in order capture real-time phenomena, such as, hurricanes, fires, floods, and tsunamis.

In the domain of *Bioinformatics*, visualization techniques are exploited in numerous tasks. For example, analyzing the large amounts of biological data produced by DNA sequencers is extremely challenging. Visual techniques can help biologist to gain insight and identify interesting “areas of genes on which to performs their experiments.

In the Big Data era, visualization techniques are extensively used in the *business intelligence* domain. *Finance markets* is one application area, where visual analytics allow to monitor markets, identify trends and perform predictions. Besides, *market research* is also an application area. Marketing agencies and in-house marketing departments analyze a plethora of diverse sources (e.g., finance data, customer behavior, social media). Visual techniques are exploited to realize task such as, identifying trends, finding emerging market opportunities, finding influential users and communities, optimizing operations (e.g., troubleshooting of products and services), business analysis and development (e.g., churn rate prediction, marketing optimization).

7 Further Reading

The literature on visualization is extensive, covering a large range of fields and many decades. Data visualization is discussed in a great number of recent introductory-level textbooks, such as [1, 62, 63, 64, 65].

Also, there are various articles discussing Big Data visualization; see [3, 4, 5, 6, 9]. Surveys of Big Data visualization systems can be found at [2, 7, 8].

In what follows we provide some surveys/studies related to issues discussed in this article:

- *Graph visualization* [66, 67, 68]
- *Hierarchical exploration* [15]
- *Visualization recommendations* [49]
- *Linked and Web data visualization* [2, 69, 70, 71]
- *High-dimensional data visualization* [72]
- *Temporal data visualization* [73]

Some of the major workshops and symposiums focusing on Big Data visualization include:

- *Workshop on Big Data Visual Exploration and Analytics* (BigVis)
- *Symposium on Big Data Visual Analytics* (BDVA)
- *Big Data Analysis and Visualization* (LDAV)

- *Workshop on Data Mining Meets Visual Analytics at Big Data era (DAVA)*
- *Workshop on Human-In-the-Loop Data Analytics (HILDA)*
- *Workshop on Data Systems for Interactive Analysis (DSIA)*
- *Workshop on Immersive Analytics: Exploring Future Interaction and Visualization Technologies for Data Analytics*

Finally, there is a great deal of information regarding visualization tools available in the Web. We mention dataviz.tools⁴ and [datavizcatalogue](http://www.datavizcatalogue.com)⁵ which are catalogs containing a large number of visualization tools, libraries and resources.

8 Cross-References

- Visualization
- Visualization Techniques
- Visualizing Semantic Data
- Graph exploration and search

References

1. M. O. Ward, G. Grinstein, and D. Keim, *Interactive Data Visualization: Foundations, Techniques, and Applications, Second Edition*. A. K. Peters, Ltd., 2015.
2. N. Bikakis and T. Sellis, “Exploration and Visualization in the Web of Big Linked Data: A Survey of the State of the Art,” in *6th Intl. Workshop on Linked Web Data Management (LWDM)*, 2016.
3. B. Shneiderman, “Extreme Visualization: Squeezing a Billion Records into a Million Pixels,” in *ACM Conference on Management of Data (SIGMOD)*, 2008.
4. J. Heer and S. Kandel, “Interactive Analysis of Big Data,” *ACM Crossroads*, vol. 19, no. 1, 2012.
5. K. Morton, M. Balazinska, D. Grossman, and J. D. Mackinlay, “Support the Data Enthusiast: Challenges for Next-generation Data-analysis Systems,” *VLDB Endowment (PVLDB)*, vol. 7, no. 6, 2014.
6. E. Wu, L. Battle, and S. R. Madden, “The Case for Data Visualization Management Systems,” *VLDB Endowment (PVLDB)*, vol. 7, no. 10, 2014.
7. P. Godfrey, J. Gryz, and P. Lasek, “Interactive Visualization of Large Data Sets,” *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 28, no. 8, 2016.
8. S. Idreos, O. Papaemmanouil, and S. Chaudhuri, “Overview of Data Exploration Techniques,” in *ACM Conference on Management of Data (SIGMOD)*, 2015.
9. J. Heer and B. Shneiderman, “Interactive Dynamics for Visual Analysis,” *Commun. ACM*, vol. 55, no. 4, 2012.
10. D. Fisher, I. O. Popov, S. M. Drucker, and M. C. Schraefel, “Trust Me, I’m Partially Right: Incremental Visualization Lets Analysts Explore Large Datasets Faster,” in *Conference on Human Factors in Computing Systems (CHI)*, 2012.

⁴ <http://dataviz.tools>

⁵ www.datavizcatalogue.com

11. Y. Park, M. J. Cafarella, and B. Mozafari, "Visualization-aware Sampling for Very Large Databases," in *IEEE Intl. Conf. on Data Engineering (ICDE)*, 2016.
12. S. Agarwal, B. Mozafari, A. Panda, H. Milner, S. Madden, and I. Stoica, "Blinkdb: Queries with Bounded Errors and Bounded Response Times on Very Large Data," in *European Conference on Computer Systems (EuroSys)*, 2013.
13. J. Im, F. G. Villegas, and M. J. McGuffin, "Visreduce: Fast and Responsive Incremental Information Visualization of Large Datasets," in *IEEE Conf. on Big Data (BigData)*, 2013.
14. L. Battle, M. Stonebraker, and R. Chang, "Dynamic Reduction of Query Result Sets for Interactive Visualizaton," in *IEEE Conf. on Big Data (BigData)*, 2013.
15. N. Elmqvist and J. Fekete, "Hierarchical Aggregation for Information Visualization: Overview, Techniques, and Design Guidelines," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 16, no. 3, 2010.
16. N. Bikakis, G. Papastefanatos, M. Skourla, and T. Sellis, "A Hierarchical Aggregation Framework for Efficient Multilevel Visual Exploration and Analysis," *Semantic Web Journal*, vol. 8, no. 1, 2017.
17. U. Jugel, Z. Jerzak, G. Hackenbroich, and V. Markl, "VDDa: Automatic Visualization-driven Data Aggregation in Relational Databases," *Journal on Very Large Data Bases (VLDBJ)*, 2015.
18. Z. Liu, B. Jiang, and J. Heer, "*imMens*: Real-time Visual Querying of Big Data," *Comput. Graph. Forum (CGF)*, vol. 32, no. 3, pp. 421–430, 2013.
19. L. D. Lins, J. T. Klosowski, and C. E. Scheidegger, "Nanocubes for Real-time Exploration of Spatiotemporal Datasets," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 19, no. 12, 2013.
20. P. Godfrey, J. Gryz, P. Lasek, and N. Razavi, "Visualization through Inductive Aggregation," in *Intl. Conf. on Extending Database Technology (EDBT)*, 2016.
21. B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations," in *IEEE Symposium on Visual Languages*, 1996.
22. J. F. R. Jr., H. Tong, J. Pan, A. J. M. Traina, C. T. Jr., and C. Faloutsos, "Large Graph Analysis in the GMine System," *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 25, no. 1, 2013.
23. J. Abello, F. van Ham, and N. Krishnan, "ASK-graphview: A Large Scale Graph Visualization System," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 12, no. 5, 2006.
24. D. Auber, "Tulip - a Huge Graph Visualization Framework," in *Graph Drawing Software*, 2004.
25. C. Tominski, J. Abello, and H. Schumann, "Cgv - an Interactive Graph Visualization System," *Computers & Graphics*, vol. 33, no. 6, 2009.
26. Z. Lin, N. Cao, H. Tong, F. Wang, U. Kang, and D. H. P. Chau, "Demonstrating Interactive Multi-resolution Large Graph Exploration," in *ICDM Workshops*, 2013.
27. M. Zinsmaier, U. Brandes, O. Deussen, and H. Strobel, "Interactive Level-of-detail Rendering of Large Graphs," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 18, no. 12, 2012.
28. D. Archambault, T. Munzner, and D. Auber, "Grouseflocks: Steerable Exploration of Graph Hierarchy Space," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 14, no. 4, 2008.
29. S. Sundara, M. Atre, V. Kolovski, S. Das, Z. Wu, E. I. Chong, and J. Srinivasan, "Visualizing Large-scale Rdf Data Using Subsets, Summaries, and Sampling in Oracle," in *IEEE Intl. Conf. on Data Engineering (ICDE)*, 2010.
30. N. Bikakis, J. Liagouris, M. Krommyda, G. Papastefanatos, and T. Sellis, "Graphvizdb: A Scalable Platform for Interactive Large Graph Visualization," in *IEEE Intl. Conf. on Data Engineering (ICDE)*, 2016.
31. E. R. Gansner, Y. Hu, S. C. North, and C. E. Scheidegger, "Multilevel Agglomerative Edge Bundling for Visualizing Large Graphs," in *IEEE Pacific Visualization Symposium (PacificVis)*, 2011.

32. O. Ersoy, C. Hurter, F. V. Paulovich, G. Cantareiro, and A. Telea, "Skeleton-based Edge Bundling for Graph Visualization," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 17, no. 12, 2011.
33. D. Phan, L. Xiao, R. B. Yeh, P. Hanrahan, and T. Winograd, "Flow Map Layout," in *IEEE Symposium on Information Visualization (InfoVis)*, 2005.
34. A. Lambert, R. Bourqui, and D. Auber, "Winding Roads: Routing Edges into Bundles," *Comput. Graph. Forum (CGF)*, vol. 29, no. 3, 2010.
35. W. Cui, H. Zhou, H. Qu, P. C. Wong, and X. Li, "Geometry-based Edge Clustering for Graph Visualization," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 14, no. 6, 2008.
36. D. Holten, "Hierarchical Edge Bundles: Visualization of Adjacency Relations in Hierarchical Data," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 12, no. 5, 2006.
37. C. D. Stolper, A. Perer, and D. Gotz, "Progressive Visual Analytics: User-driven Visual Exploration of In-progress Analytics," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 20, no. 12, 2014.
38. A. Kalinin, U. Çetintemel, and S. B. Zdonik, "Interactive Data Exploration Using Semantic Windows," in *ACM Conference on Management of Data (SIGMOD)*, 2014.
39. I. Alagiannis, R. Borovica, M. Branco, S. Idreos, and A. Ailamaki, "Nodb: Efficient Query Execution on Raw Data Files," in *ACM Conference on Management of Data (SIGMOD)*, 2012.
40. M. Olma, M. Karpathiotakis, I. Alagiannis, M. Athanassoulis, and A. Ailamaki, "Slalom: Coasting through Raw Data Via Adaptive Partitioning and Indexing," *VLDB Endowment (PVLDB)*, vol. 10, no. 10, 2017.
41. K. Zoumpatianos, S. Idreos, and T. Palpanas, "Indexing for Interactive Exploration of Big Data Series," in *ACM Conference on Management of Data (SIGMOD)*, 2014.
42. M. Karpathiotakis, M. Branco, I. Alagiannis, and A. Ailamaki, "Adaptive Query Processing on Raw Data," *VLDB Endowment (PVLDB)*, vol. 7, no. 12, 2014.
43. Y. Tian, I. Alagiannis, E. Liarou, A. Ailamaki, P. Michiardi, and M. Vukolic, "Dinodb: An Interactive-speed Query Engine for Ad-hoc Queries on Temporary Data," *IEEE Trans. Big Data*, vol. 3, no. 3, 2017.
44. L. Battle, R. Chang, and M. Stonebraker, "Dynamic Prefetching of Data Tiles for Interactive Visualization," in *ACM Conference on Management of Data (SIGMOD)*, 2016.
45. F. Tauheed, T. Heinis, F. Schürmann, H. Markram, and A. Ailamaki, "SCOUT: Prefetching for Latent Feature Following Queries," *VLDB Endowment (PVLDB)*, vol. 5, no. 11, 2012.
46. P. Jayachandran, K. Tunga, N. Kamat, and A. Nandi, "Combining User Interaction, Speculative Query Execution and Sampling in the Dice System," *VLDB Endowment (PVLDB)*, vol. 7, no. 13, 2014.
47. S. Chan, L. Xiao, J. Gerth, and P. Hanrahan, "Maintaining Interactivity While Exploring Massive Time Series," in *IEEE VAST*, 2008.
48. H. A. Khan, M. A. Sharaf, and A. Albarrak, "Divide: Efficient Diversification for Interactive Data Exploration," in *Intl. Conf. on Scientific and Statistical Database Management (SSDBM)*, 2014.
49. M. Vartak, S. Huang, T. Siddiqui, S. Madden, and A. G. Parameswaran, "Towards Visualization Recommendation Systems," *SIGMOD Record*, vol. 45, no. 4, 2016.
50. A. Key, B. Howe, D. Perry, and C. R. Aragon, "Vizdeck: Self-organizing Dashboards for Visual Analytics," in *ACM Conference on Management of Data (SIGMOD)*, 2012.
51. H. Ehsan, M. A. Sharaf, and P. K. Chrysanthis, "Muve: Efficient Multi-objective View Recommendation for Visual Data Exploration," in *IEEE Intl. Conf. on Data Engineering (ICDE)*, 2016.
52. J. D. Mackinlay, P. Hanrahan, and C. Stolte, "Show Me: Automatic Presentation for Visual Analysis," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 13, no. 6, 2007.
53. G. A. Atemezeng and R. Troncy, "Towards a Linked-data Based Visualization Wizard," in *Workshop on Consuming Linked Data (COLD)*, 2014.

54. K. Thellmann, M. Galkin, F. Orlandi, and S. Auer, "Linkdaviz - Automatic Binding of Linked Data to Visualizations," in *Intl. Semantic Web Conf. (ISWC)*, 2015.
55. M. Vartak, S. Madden, A. G. Parameswaran, and N. Polyzotis, "SEEDB: Automatically Generating Query Visualizations," *VLDB Endowment (PVLDB)*, vol. 7, no. 13, 2014.
56. K. Wongsuphasawat, D. Moritz, A. Anand, J. D. Mackinlay, B. Howe, and J. Heer, "Voyager: Exploratory Analysis Via Faceted Browsing of Visualization Recommendations," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 22, no. 1, 2016.
57. G. Wills and L. Wilkinson, "Autovis: Automatic Visualization," *Information Visualization*, vol. 9, no. 1, 2010.
58. B. Mutlu, E. E. Veas, and C. Trattner, "Vizrec: Recommending Personalized Visualizations," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 6, no. 4, 2016.
59. D. Gotz and Z. Wen, "Behavior-driven Visualization Recommendation," in *Intl. Conference on Intelligent User Interfaces (IUI)*, 2009.
60. S. Kandel, R. Parikh, A. Paepcke, J. M. Hellerstein, and J. Heer, "Profiler: Integrated Statistical Analysis and Visualization for Data Quality Assessment," in *Intl. Working Conference on Advanced Visual Interfaces (AVI)*, 2012.
61. E. Wu and S. Madden, "Scorpion: Explaining Away Outliers in Aggregate Queries," *VLDB Endowment (PVLDB)*, vol. 6, no. 8, 2013.
62. S. Murray, *Interactive Data Visualization for the Web*. O'Reilly Media, Inc., 2013.
63. M. O. Ward, G. G. Grinstein, and D. A. Keim, *Interactive Data Visualization - Foundations, Techniques, and Applications*. A K Peters, 2010.
64. D. A. Keim, J. Kohlhammer, G. P. Ellis, and F. Mansmann, *Mastering the Information Age - Solving Problems with Visual Analytics*. Eurographics Association, 2010.
65. C. Johnson and C. Hansen, *Visualization Handbook*. Academic Press, Inc., 2004.
66. T. von Landesberger, A. Kuijper, T. Schreck, J. Kohlhammer, J. J. van Wijk, J. Fekete, and D. W. Fellner, "Visual Analysis of Large Graphs: State-of-the-art and Future Research Challenges," *Comput. Graph. Forum (CGF)*, vol. 30, no. 6, 2011.
67. I. Herman, G. Melançon, and M. S. Marshall, "Graph Visualization and Navigation in Information Visualization: A Survey," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 6, no. 1, 2000.
68. G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall, 1999.
69. A. Dadzie and E. Pietriga, "Visualisation of Linked Data - Reprise," *Semantic Web*, vol. 8, no. 1, 2017.
70. N. Marie and F. L. Gandon, "Survey of Linked Data Based Exploration Systems," in *Intl. Workshop on Intelligent Exploration of Semantic Data*, 2014.
71. A. Dadzie and M. Rowe, "Approaches to Visualising Linked Data: A Survey," *Semantic Web*, vol. 2, no. 2, 2011.
72. S. Liu, D. Maljovec, B. Wang, P. Bremer, and V. Pascucci, "Visualizing High-dimensional Data: Advances in the Past Decade," *IEEE Trans. Vis. Comput. Graph. (TVCG)*, vol. 23, no. 3, 2017.
73. E. Wohlfart, W. Aigner, A. Bertone, and S. Miksch, "Comparing Information Visualization Tools Focusing on the Temporal Dimensions," in *Intl. Conference on Information Visualisation (IV)*, 2008.