

Towards Big Data Interactive Visualization in Ambient Intelligence Environments

Giannis Drossis¹, George Margetis^{1(✉)}, and Constantine Stephanidis^{1,2}

¹ Institute of Computer Science, Foundation for Research and Technology - Hellas (FORTH),
N. Plastira 100, Vassilika Vouton, 70013 Heraklion, Crete, Greece
{drossis, gmarget, cs}@ics.forth.gr

² Computer Science Department, University of Crete, Heraklion, Crete, Greece

Abstract. Big Data visualization relies on an interdisciplinary research area that includes mass data storage and retrieval, operations, analytics, security, ethics as well as visualization and interaction with end users. This paper reports on the characteristics of Big Data systems, mainly focusing on information visualization and discusses a number of methods towards this direction, analyzing research issues and challenges that emerge. Additionally, this paper discusses new approaches for Big Data visualization in the context of Ambient Intelligence (AmI) environments, highlighting new aspects in the field in respect to information presentation and natural user interaction. Furthermore, a scenario of Big Data visualization in AmI environments is presented, aiming at bringing to surface the new potential of such approaches in terms of interaction simplification, and adaptation to the context of use.

Keywords: Big Data · Big Data visualization · Big Data interaction · Ambient Intelligence · Data centre infrastructure management

1 Introduction

The term “Big Data” emerged as a scientific domain during the 1990s and gradually became a field of major academic interest during the 2000s. Since 2010, Big Data became a focus of attention across various sectors, including trade, business intelligence, large-scale organizations and start-up companies.

A widely accepted definition of Big Data is still under discussion and there is little progress regarding a commonly acceptable answer to the fundamental question of how big the data has to be to qualify as ‘big data’. De Mauro et al. [1] combine several definitions for Big Data as “representing the Information assets characterized by such a High Volume, Velocity and Variety to require specific Technology and Analytical Methods for its transformation into Value”.

Laney [21] introduced one of the first and most widely embraced attempts to define Big Data. The author suggested that Volume, Variety, and Velocity (or else the three Vs) are the three dimensions of challenges in data management. Volume refers to the magnitude of data. The thresholds that discriminate data from big data however are vague, as they keep increasing due to the soaring of sensors supplying data. Variety refers to the

heterogeneity of the data, while velocity refers to the rate at which data is generated, retrieved and processed. Data creation rate is currently increasing rapidly with smart-phones and sensors, and is expected to multiply in the era of the Internet of Things.

In addition to the three Vs, other dimensions that characterize other aspects of Big Data are described in literature. The factor of veracity [12], which describes the quality of the data, is an additional important characteristic that provides assessment of the data. Finally, another term is value [26], which describes the dimension of the data values in the context of Big Data integration.

Due to the vague nature of Big Data, several research fields are involved in the process of its management and manipulation. Such fields include, but are not limited to, storage, operations, analytics, security, ethics and visualization. This paper will focus on the classification of the major Big Data types, on the existing visualization techniques and on potential future directions of visualization and interaction with Big Data in the context of Ambient Intelligence environments.

2 Big Data Analytics

Big Data Analytics refer to the process of studying, exploring and reporting on Big Data, aiming at assisting the human decision making process. Apart from reporting and visualization, operations that belong to several research fields are involved, such as predictive analytics, data mining, anomaly detection, statistical analysis and text analytics. As mentioned in [10], research challenges regarding Big Data analytics involve both structural approaches and visualization concerns. The structural approaches refer to the infrastructure that allows making operations on Big Data, while visualization concerns encompass the nature of the data source in terms of presentation.

A fundamental characteristic of Big Data is its multidimensional nature. Big Data comprises massive information with several aspects. For instance, information about a specific car not only contains data about the car itself, such as model, production start date, horsepower, acceleration, but also data regarding sales: how many cars were sold, the sale location, the color of the cars, etc. This information comprises different dimensions of the data that create different perspectives, which can be retrieved and analyzed on demand.

2.1 Analytical Processing and Dimensions of Big Data

The first step towards visualizing multidimensional data sets includes the identification of the involved aspects that need to be presented. Current approaches that cope with this issue mainly involve OLAP [5] (Online Analytical Processing), which aims to provide a mechanism for the analysis of Big Data from multiple different perspectives. OLAP is applied as a means of interactive filtering out extrinsic information. Several data storage models for the core implementation of OLAP are found in literature, including Multidimensional OLAP systems (MOLAP), Relational OLAP (ROLAP) and Hybrid OLAP (HOLAP). This paper will not focus on the different core approaches that allow OLAP operations, but instead on the visualization aspects that arise at a higher level.

2.2 Temporal Big Data

Time classification is an aspect of Big Data, which is frequently meaningful. It constitutes a dimension that is often stored, either as supplementary information or as major descriptive value. Users are accustomed to perceiving time as an additional dimension of Big Data and are therefore able to easily perceive it as a dimension to any type of information. Time constitutes an important factor in various contexts, including decision-making [7].

2.3 Geospatial Big Data

Another type of information which is often an integral part of Big Data is space. The datasets can be static or dynamic in terms of location. Dynamic Big Data can be acquired through the usage of location-aware devices and the adoption of Geographical Information Systems (GIS) and are correlated both to space and time. The quality of data, which imposes an important issue in terms of geospatial data [22], is expected to improve especially in the context of the Internet of Things.

Li et al. [23] use Big Data 3D visualization in order to display infrastructure in Shenzhen and take geographic statistical analysis in consideration for assisting data analysis and enhancing the decision making process for social service agencies. Big Data in the context of urban environments in general correlate data sets with both time [18] and space [3].

3 Big Data Visualization Techniques

3.1 OLAP Visualization

The OLAP tools enable users to analyze multidimensional data interactively from various perspectives. OLAP consists of five basic analytical operations: consolidation (roll-up), drill-down, pivoting, slicing and dicing [14]. The applied operations change the selection of the visualized components on the fly and update the view according to the user's actions.

A traditional interface for analyzing OLAP data is a pivot table, or cross-tab, which is a multidimensional spreadsheet produced by specifying one or more measures of interest and selecting dimensions to serve as vertical (and, optionally, horizontal) axes for summarizing the measures [9]. Pivot tables are a widespread visualization model, which provide a detailed data presentation to users that are familiar with it. Additionally, they can be easily transformed to serve additional dimensions of information through the application of color-coding or rows/columns merging. However, the tables' efficiency declines for larger data sets, as users are unable to locate specific information, recognize patterns or get an overview of the displayed data sets (Fig. 1).

Another useful visualization component are parallel coordinates [15, 19], which allow the display of multiple data dimensions on a 2D plane. The concept of parallel coordinates is the concurrent visualization of different values in a row, one after another. For instance, in the case of car comparison, car values regarding horsepower, fuel

Category	Type	2010	2011	2012	2013	Total
Fruit	Apples	551	477	604	648	2280
	Bananas	988	763	893	825	3469
	Oranges	226	261	219	254	960
	<i>Subtotal</i>	<i>1765</i>	<i>1501</i>	<i>1716</i>	<i>1727</i>	
Vegetables	Lettuce	423	461	489	502	1875
	Tomatoes	335	356	376	328	1395
	<i>Subtotal</i>	<i>758</i>	<i>817</i>	<i>865</i>	<i>830</i>	
Total		2523	2318	2581	2557	

Fig. 1. An example pivot table

consumption, car dimensions, acceleration, etc., would be meaningful to be displayed on the same plot. This visualization technique holds the advantage of correlation between different data dimensions, e.g., horsepower and acceleration, which could depict the impact of one dimension on another. However, this applies only in the case where the dimensions are rendered adjacently; if another dimension intervenes between them, the visual clue disappears and the correlation is impossible to discover. Additionally, parallel coordinates are unable to visualize non-numerical values and can be difficult to comprehend for non-expert users (Fig. 2).

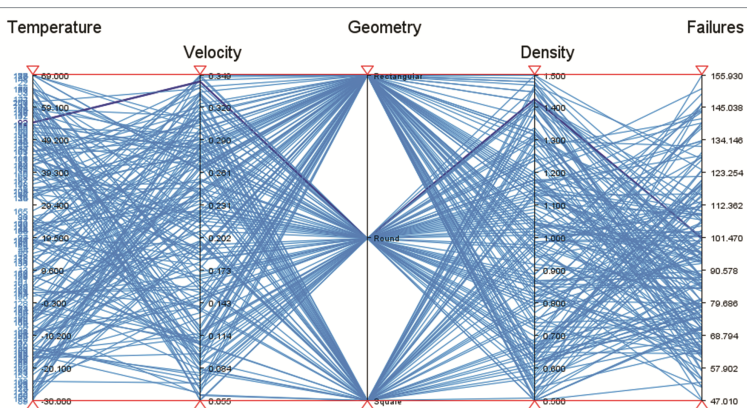


Fig. 2. An indicative example of parallel coordinates created using XDAT [37]

In order to address the aforementioned drawbacks of tables and parallel coordinates, techniques such as pie charts, plots, graphs and trees are employed by state of the art commercial tools. The leading software for Big Data Visualization, Business Intelligence and Analytics is Tableau [24], while other approaches include qlikview [29], Microsoft's Power BI [28] and others.

Tableau [24] provides default tools for rendering each specific data type, which can be overridden by the user on demand. Each visualization is selected to optimally represent a specific data type: for instance, cross-tabs are used for the visualization of discrete categorical data and lines for continuous quantitative information.

Even though OLAP tools are very efficient for Big Data retrieval and on demand visualization, they usually lack exploratory functionality. Furthermore, due to the complex nature of the operations supplied, OLAP tools tend to be cumbersome for novice users to manipulate.

3.2 Big Data 3D Visualization

Traditional visualization techniques fall short in terms of efficient and intuitive display of the corresponding data sets; therefore, the need for a rich interactive visualization still constitutes both a business and a research challenge.

In order to fill in the gap of inefficient visualization, some approaches in literature involve 3D visualization for OLAP [2, 20]. However, 3D visualization approaches are not yet very popular.

The inclusion of an additional dimension to the visualization adds up to the efficiency of depicting the dimensions of Big Data. Since the human brain is trained to sense and act in three dimensions, it is optimized to perceive three dimensions in a natural manner. The third dimension constitutes an aspect of the rendered information easily perceivable by the user, thus enhancing potential exploration in a virtual world.

On the other hand, 3D visualizations have certain drawbacks. In general, they present a steep learning curve, as inexperienced users usually tend to be unable to orientate themselves in the virtual three-dimensional space or to manipulate the interface. The lack of perception of 3D space does not only involve the virtual space itself, but also the lack of the exploration and identification of the visualized information. Interaction complexity with 3D user interfaces is also a significant burden due to the additional degrees of freedom, requiring complex manipulation controls and a rich interaction vocabulary.

3.3 Virtual and Augmented Reality

Virtual and augmented reality environments form an emerging approach that is capable of providing Big Data visualizations. Such environments in the context of Big Data constitute interdisciplinary efforts, combining the areas of 3D graphics, stereoscopic environments, computer vision and Big Data querying. The main advantage of virtual and augmented reality environments is better user experience and immersion, which allows the better perception of the visualized geometry. Furthermore, in comparison to traditional 3D visualizations, the users perceive themselves in the context of the visualization and thus orientate themselves more easily.

Helbig et al. [17] use a virtual reality environment to visualize massive data in the context of Weather Research Forecast. Another interesting approach is the immersive visualization of a landscape in Mars [11], augmented with data describing the surface characteristics.

Future challenges on applying virtual and augmented reality to Big Data visualizations include multimodal interaction, display and equipment limitations [25]. Virtual and augmented reality is a growing research field, mainly due to the emergence of devices like the Microsoft Kinect and LeapMotion that provide more natural interaction based on gestures and Oculus Rift, which puts virtual reality in play again in terms of mainstream

visualization technique. However, still several challenges exist towards incorporating and enhancing traditional 2D desktop approaches in virtual space as well as developing a suitable infrastructure at the side of Big Data to support additional needs that may rise.

3.4 Graph Visualization

Graphs are a common technique used for displaying the correlation between different entities. Their main advantage is the user's ability of starting from a specific node and exploring neighboring nodes, especially when visualizing data sets that describe networks or relationships. The survey reported by Beck et al. [4] describes a trend towards the combination of graphs with interactive timelines in order to include potential temporal characteristics of the information.

Furthermore, graph visualizations simplify exploration by providing operations like sampling, filtering, partitioning and clustering [27], while they can also support several abstraction layers [6] in order to provide meaningful views according to the scope of the visualization, ranging from overview to detailed view.

Even though graphs can be very helpful in illustrating specific aspects of Big Data, they tend to focus only on one aspect of the data, which is the interconnection between the various nodes. Moreover, graphs are meaningful only if the data that they present are coherent, and are not suitable to illustrate other aspects such as comparison between data and temporal relationships.

3.5 Exploratory Data Analysis

Exploratory Analytics, or else Discovery Analytics, refers to the process of using visualization exploration techniques, in the context of Big Data, which aims at discovering new facts or characteristics of Big Data that users were previously unaware of [31, 34, 36]. Heer and Shneiderman [16] present a widely adopted taxonomy for interactive dynamics regarding visual analysis. The proposed taxonomy groups tasks in three high level categories: data and view specification, view manipulation and analysis process and provenance.

Faceted navigation, also mentioned as data and view specification [16], refers to the process of applying specific filters to the data sets provided, in order to focus on the subset of interest. Faceted navigation can combine multiple visualization techniques to apply the most suitable ones to the corresponding type. Such an example is EDEN [33], where the authors use parallel coordinates and geographic visualizations to interactively refine the displayed values and thus offer exploratory analysis of Big Data by exploring relationships between entities. Another example of interactive faceted exploration is discussed in [36], which combines automatically generated and manually specified visualizations in order to improve support for data exploration.

4 Comprehension of Big Data

What still proves to be a difficult task in the area of Big Data is offering the possibility to the users to comprehend the dimensions of Big Data. Even though much work focuses on analyzing the data itself by making operations in order to view and examine specific

aspects of the data, less effort has focused on creating the optimal environments for the depiction of the available aspects and offering exploration capabilities.

Comprehending Big Data involves two different roles: (a) experts who are either familiar with the visualized data or at least are accustomed to the concepts of Big Data (e.g., data scientists and analysts), and (b) inexperienced end users, who are unaccustomed to perceiving and manipulating Big Data. The visualization approach for these different roles varies significantly, as different expectations are raised in the corresponding cases.

The vast majority of the existing Big Data visualization techniques focus on the experts. On the other hand, occasional or inexperienced users should be provided with a more user-friendly, exploratory view that allows browsing without requiring knowledge of the existing data operators or the visualized data domain.

5 Displays of Big Data Visualizations and User Collaboration

Augmented and virtual reality techniques require specialized displays, while the vast majority of other visualizations relies on desktop displays. On the other hand, the interest of non-expert end users is expected to increase, requiring access to Big Data in a user-friendly and easily perceived manner. Therefore, a forthcoming issue will be the access to Big Data visualizations for a broad audience through different platforms and devices, such as mobile devices [38].

Furthermore, multi-device environments with shared displays and mobile devices [30] will support teamwork and collaborative exploration of visualizations. To this direction, Donalek et al. [11] discuss means of sharing displays between users.

6 Towards Ambient Intelligence

Ambient Intelligence envisions a future where technology is interweaved with everyday living environments, anticipates users' needs, and provides natural interaction with digital information [35].

In respect to human-computer interaction, Ambient Intelligence (AmI) environments involve primarily two aspects: context awareness and natural interfaces [8]. Context awareness includes the use of emerging technologies to infer the context of the interaction (e.g., the location and the activities of the user), whereas natural interfaces refer to human communication capabilities and implicit actions that should be employed as a means of interaction in AmI environments, instead of the explicit input used in traditional human-computer interaction.

Ambient Intelligence environments have the potential to support the visualization of Big Data and its multiple dimensions due to the wide range of displays and interactions combined, which can be both implicit and explicit [32].

6.1 The AmI Data Centre Use Case

Data centre infrastructure management constitutes a Big Data use case, as massive data is generated and collected in real time. Information such as the servers' state, network traffic and temperature is continuously created along with logical and physical dependencies that seldom vary.

The massive growth of data volume, following Moore's law, results into the need for further storage space. Consequently, the existing data centres are expected to grow both in size and in count; thus, their monitoring, maintenance and management involves comprehensive environment sensing, where the environment is the data centre itself.

Data centres consist of multiple server racks, which in turn contain several servers. Each server holds a diversity of metrics that are required for its maintenance and optimization, including anomaly detection, workload balance and power consumption. The servers' values are observed and displayed by real time visualization tools, such as the one presented by Drossis et al. [13].

The integration of Ambient Intelligence could enhance infrastructure management in various perspectives. The data centre control room can be greatly improved by coalescing existing management applications with sensors and pervasive displays, that would allow sensing in the area and act in combination with activities of the employees and the state of the data centre itself. Additionally, navigation and the finding of specific components is another field that requires user assistance. The following use case scenario depicts some aspects that involve Big Data visualization in the context of Ambient Intelligence.

6.2 Use Case Scenario - A Day in the Data Centre

A 3D visualization of the data centre is displayed at the largest display of the AmI data centre's control room, showing the current state of all the servers. The traffic between the connections is monitored in real-time and overlaid in the visualized servers' layout to the corresponding responsible employee. The employee spots a potential issue and shares his view with his colleagues, having the selected connections between servers shown in the large display. The real-time analysis of the data center performance, by the AmI environment, discovers a bottleneck in traffic, and raises a warning to the personnel, providing a historical view of the traffic for the specific data centre to a side large display. After the analysis of the traffic history, the data centre's personnel decides to install additional servers to improve the performance.

When the new servers are delivered, an employee places them on the smart working bench, which is capable of recognizing the servers' configuration by reading their serial numbers. The smart bench suggests, based on the current traffic demand, alternative plans of the new servers' placement, illustrated on a floorplan projected next to them, with the compatible empty slots highlighted. The administrator of the data centre selects the preferred position among the available options, using a hand gesture, notifying this way that the hardware is ready to be installed. Upon installation, all the users are notified about the addition and the visualizations of the data centre are updated accordingly.

Afterwards, a component with very high temperature is identified by automated anomaly detection mechanisms. The environment triggers an alarm notifying the responsible employee, who in turn filters the visualization to display only the servers with high temperatures to verify that the reporting sensor is not problematic. Furthermore, the details of the anomaly are shown on demand by bringing to front the temperature control, which displays the previous values along with the expected ones.

Upon user intervention, the alarm can be acknowledged by the user or examined in detail. Since the anomaly requires on spot actions, the employee uses his own smartphone to view the path towards the faulty component, which is highlighted. Upon reaching the defined rack, he uses an augmented reality view to find the detected server slot in the rack with the reported temperature. The high temperature was reported due to a sensor malfunction and therefore the employee replaces the faulty component and marks the issue as resolved.

6.3 Big Data and Ambient Intelligence

The scenario presented in the previous section illustrates the benefits that Ambient Intelligence offers for Big Data centre monitoring and maintenance. In more details, the control room of such a data centre provides visualization facilities that are able to display thousands of servers according to the context of use (the large display shows generic information, whereas the users see a personalized view). The visualization displays relationships between servers (network traffic) and past values (temporal information), and users are able to collaborate (employee's view is shared).

The smart bench is able to identify physical objects (new servers) but also to suggest potential installation locations (specifications of the servers are known and consequently the environment acts in a smart manner). The installation is sensed by the environment and the visualization is updated immediately (real time visualization).

The high temperature alarm is personalized and raised to the responsible employee for action. The employee applies filtering (faceted navigation) and browses towards the desired area (exploration). Additionally, the temperature control is specialized to suit the needs of its content (data optimal visualization). Finally, the employee uses his own smartphone to find the area of interest in a natural manner (using navigation with augmented reality).

7 Conclusion

Big Data visualizations constitute an interdisciplinary area of research that combines different technologies. This paper has discussed the characteristics of Big Data, the existing visualization techniques and the potential benefits of Ambient Intelligence for Big Data visualization. A use case scenario of a data centre was described, illustrating the way that Ambient Intelligence can enhance Big Data visualizations in the context of data centre infrastructure management.

Ambient Intelligence has the potential to fill in gaps of existing Big Data visualization techniques by incorporating the context of use to the visualization and supplying rich multimodal interaction. Ambient Intelligence can reduce the dimensions of Big Data and adapt the display of the data rendered in a smart and natural manner, thus reducing the interface complexity and the user cognitive load. Moreover, Ambient Intelligence can support and promote collaborative work and real-time human-human interaction.

Acknowledgements. The work reported in this paper has been conducted in the context of the AmI Programme of the Institute of Computer Science of the Foundation for Research and Technology-Hellas (FORTH).

References

1. De Mauro, A., Greco, M., Grimaldi, M.: What is big data? A consensual definition and a review of key research topics. In: AIP Conference Proceedings, vol. 1644, no. 1 (2015)
2. Ammoura, A., Zai'ane, O.R., Göbel, R.: Towards a novel OLAP interface for distributed data warehouses. In: Kambayashi, Y., Winiwarter, W., Arikawa, M. (eds.) DaWaK 2001. LNCS, vol. 2114, pp. 174–185. Springer, Heidelberg (2001)
3. Azri, S., Ujang, U., Castro, F.A., Rahman, A.A., Mioc, D.: Classified and clustered data constellation: an efficient approach of 3D urban data management. ISPRS J. Photogramm. Remote Sens. **113**, 30–42 (2016)
4. Beck, F., Burch, M., Diehl, S., Weiskopf, D.: The state of the art in visualizing dynamic graphs. EuroVis STAR (2014)
5. Berson, A., Smith, S.J.: Data Warehousing, Data Mining, and OLAP. McGraw-Hill Inc., New York (1997)
6. Bikakis, N., Liagouris, J., Krommyda, M., Papastefanatos, G.: graphVizdb: A Scalable Platform for Interactive Large Graph Visualization
7. Chandramouli, B., Goldstein, J., Duan, S.: Temporal analytics on big data for web advertising. In: 2012 IEEE 28th International Conference on Data Engineering (ICDE). IEEE (2012)
8. Cook, D.J., Augusto, J.C., Jakkula, V.R.: Ambient intelligence: technologies, applications, and opportunities. Pervasive Mob. Comput. **5**(4), 277–298 (2009)
9. Cuzzocrea, A., Mansmann, S.: OLAP visualization: models, issues, and techniques. In: Encyclopedia of Data Warehousing and Mining, pp. 1439–1446 (2009)
10. Cuzzocrea, A., Song I.-Y., Davis, K.C.: Analytics over large-scale multidimensional data: the big data revolution! In: Proceedings of the ACM 14th International Workshop on Data Warehousing and OLAP. ACM (2011)
11. Donalek, C., Djorgovski, S.G., Cioc, A., Wang, A., Zhang, J., Lawler, E., Yeh, S., Mahabal, A., Graham, M., Drake, A., Davidoff, S., Norris, J.S., Longo, G.: Immersive and collaborative data visualization using virtual reality platforms. In: 2014 IEEE International Conference on Big Data (Big Data), pp. 609–614. IEEE, October 2014
12. Dong, X.L., Srivastava, D.: Big data integration. In: 2013 IEEE 29th International Conference on Data Engineering (ICDE). IEEE (2013)
13. Drossis, G., Birliraki, C., Patsiouras, N., Margetis, G., Stephanidis, C.: 3-D visualization of large-scale data centres. In: Cloud Computing and Services Science. Springer International Publishing, New York (2016)
14. Gray, J., Chaudhuri, S., Bosworth, A., Layman, A., Reichart, D., Venkatrao, M., Pellow, F., Pirahesh, H.: Data cube: a relational aggregation operator generalizing group-by, cross-tab, and sub-totals. Data Min. Knowl. Discov. **1**(1), 29–53 (1997)
15. Hauser, H., Ledermann, F., Doleisch, H.: Angular brushing of extended parallel coordinates. In: IEEE Symposium on Information Visualization. INFOVIS 2002. IEEE (2002)
16. Heer, J., Shneiderman, B.: Interactive dynamics for visual analysis. Queue **10**(2), 30 (2012)
17. Helbig, C., Bauer, H.S., Rink, K., Wulfmeyer, V., Frank, M., Kolditz, O.: Concept and workflow for 3D visualization of atmospheric data in a virtual reality environment for analytical approaches. Environ. Earth Sci. **72**(10), 3767–3780 (2014)

18. Hoppenbrouwer, E., Louw, E.: Mixed-use development: theory and practice in Amsterdam's Eastern Docklands. *Eur. Plann. Stud.* **13**(7), 967–983 (2005)
19. Inselberg, A.: The plane with parallel coordinates. *Vis. Comput.* **1**(2), 69–91 (1985)
20. Lafon, S., Bouali, F., Guinot, C., Venturini, G.: On studying a 3D user interface for OLAP. *Data Min. Knowl. Discov.* **27**(1), 4–21 (2013)
21. Laney, D.: 3D data management: controlling data volume, velocity and variety. *META Gr. Res. Note* **6**, 70 (2001)
22. Li, S., Dragicevic, S., Castro, F.A., Sester, M., Winter, S., Coltekin, A., Pettit, C., Jiang, B., Haworth J., Stein A., Cheng, T.: Geospatial big data handling theory and methods: a review and research challenges. *ISPRS J. Photogramm. Remote Sens.* (2015)
23. Li, X., Lv, Z., Zhang, B., Wang, W., Feng, S., Hu, J.: WebVRGIS based city bigdata 3D visualization and analysis (2015). arXiv preprint [arXiv:1504.01051](https://arxiv.org/abs/1504.01051)
24. Mackinlay, J.D., Hanrahan, P., Stolte, C.: Show me: automatic presentation for visual analysis. *IEEE Trans. Vis. Comput. Graph.* **13**(6), 1137–1144 (2007). <http://www.tableau.com/>
25. Olshannikova, E., Ometov, A., Koucheryavy, Y., Olsson, T.: Visualizing big data with augmented and virtual reality: challenges and research agenda. *J. Big Data* **2**(1), 1–27 (2015)
26. Pascal, H., Janowicz, K.: Linked data, big data, and the 4th paradigm. *Semant. Web* **4**(3), 233–235 (2013)
27. Pienta, R., Abello, J., Kahng, M., Chau, D.H.: Scalable graph exploration and visualization: sensemaking challenges and opportunities. In: 2015 International Conference on Big Data and Smart Computing (BigComp). IEEE (2015)
28. PowerBI: powerbi.microsoft.com. Accessed 10 Mar 2016
29. Qlik: www.qlik.com. Accessed 10 Mar 2016
30. Roberts, J.C., Ritsos, P.D., Badam, S.K., Brodbeck, D., Kennedy, J., Elmqvist, N.: Visualization beyond the desktop—the next big thing. *Comput. Graph. Appl. IEEE* **34**(6), 26–34 (2014)
31. Russom, P.: Big data analytics. TDWI Best Practices Report, Fourth Quarter, pp. 1–35 (2011)
32. Schmidt, A.: Interactive context-aware systems interacting with ambient intelligence. In: *Ambient Intelligence*, p. 159 (2005)
33. Steed, C.A., Ricciuto, D.M., Shipman, G., Smith, B., Thornton, P.E., Wang, D., Shi, X., Williams, D.N.: Big data visual analytics for exploratory earth system simulation analysis. *Comput. Geosci.* **61**, 71–82 (2013)
34. Tukey, J.W.: *Exploratory Data Analysis*, pp. 2–3. Addison-Wesley, Reading (1977)
35. Weber, W., Rabaey, J., Aarts, E.H.L.: *Ambient Intelligence*. Springer Science & Business Media, Berlin (2005)
36. Wongsuphasawat, K., Moritz, D., Anand, A., Mackinlay, J.: Voyager: exploratory analysis via faceted browsing of visualization recommendations. *IEEE Trans. Vis. Comput. Graph.* **22**(1), 649–658 (2016)
37. XDAT: X-dimensional Data Analysis Tool. www.xdat.org. Accessed 10 Mar 2016
38. Zissis, D., Lekkas, D., Koutsabasis, P.: Design and development guidelines for real-time, geospatial mobile applications: lessons from 'MarineTraffic'. In: Daniel, F., Papadopoulos, G.A., Thiran, P. (eds.) *MobiWIS 2013*. LNCS, vol. 8093, pp. 107–120. Springer, Heidelberg (2013)