Kicking Analysts Out of the Meeting Room: Supporting Future Data-driven Decision Making with Intelligent Interactive Visualization Systems

Yi Han*

National Sun Yat-sen University

Abstract

Today's data-driven decisions are largely dependent on professional analysts conducting analysis and generating visualizations for decision makers. These middlemen between data and decision makers may induce cost and trust issues in the generated visualizations. To overcome these issues, I envision a future scenario where intelligent interactive visualization systems may replace analysts in the decision-making process when the analyses and visualizations are relatively simple. However, three gaps need to be addressed before the future scenario could be realized. In this paper, I will discuss these gaps, propose potential solutions, and hope to raise a discussion on the future role of visualization systems for data-driven decision-making.

Keywords: Interactive Visualization, Business Intelligence, Data-driven Decisions.

Index Terms: Human-centered computing—Visualization—Visualization application domains—Visual analytics

1 Introduction

Compared to intuition-driven decisions, data-driven decisions are made with more confidence and could lead to better results [2, 30]. With the arrival of the big data era, huge amounts of data are being collected, analyzed, and often visualized to help decision makers quickly understand them for data-driven decisions. The visualized data can also enhance communication among decision makers and other stakeholders who may be impacted by the decisions [7].

Currently, the analysis and visualization of data depends on professional analysts, who act as the middlemen between data and decision makers. These analysts are essential today because most decision makers cannot conduct analysis or generate visualizations on their own. A typical data-driven decision-making process is illustrated in Figure 1. The process begins with a decision maker encountering a business problem that needs a data-driven decision. Therefore, a request for data is sent to the analysts, who then find the data, generate visualizations, and make a story to report back their findings. This report typically happens in a decision-making meeting where the decision maker, sometimes with the help of a group of consultants, processes the information, considers options, and makes a decision [14, 16].

However, having a middleman induces costs and requires trust. What if an analyst chooses to tune some parameters or selectively show data that favor a certain decision? The decision maker and the consultants, hereon referred to as the decision team, will have to trust the analysts, the data, and the analysis process, in order to trust the generated visualizations and stories for making a confident data-driven decision. This distrust is especially likely to happen when the visualized data contradict with the decision team's intuition. With the increasing volume and complexity in the data, this trust issue will only become higher in the future.

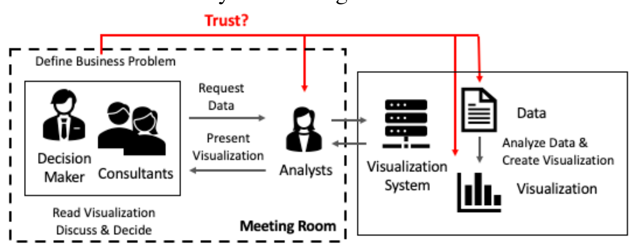

Figure 1: Illustration of a typical data-driven decision-making process. Trust issues emerge in various places.

Three types of trust issues could emerge in this decision-making process. The first type is from the lack of knowledge about the data. Research showed that insufficient disclosure of the underlying data can cause distrust in the generated visualization [8]. For example, if a decision team does not know which data are collected, then they will not know if all the necessary information is properly visualized. The second type is from the lack of visibility in the data processing, analysis, and visualization processes of the analysts. Research showed that the more the underlying data are processed, the less they are perceived as trustworthy [17]. A lack of disclosure of the analysts' decisions on how they analyzed and generated visualizations can also cause distrust in the visualizations [8]. For example, without knowing how 2D projections of topic modeling results are created, readers may misinterpret the relations among topics due to visual artifacts [4]. The third type is from the lack of familiarity with the visualization. When a visualization is less familiar to the users, it will be trusted less [5]. For example, when presented with unfamiliar advanced interactive visualizations, decision makers are more likely to dismiss them due to the increased cognitive load to understand them [5, 22]. These trust issues could all be caused by data analysts being the middlemen in the current decision-making process (Figure 1).

2 Overcoming the Trust Issue: Cut out the Middleman

In human history, there has been a trend in using technology to cut out the middleman. For example, in the past, an accountant was needed to file tax returns in the US. But today self-service tax filing tools (e.g., TurboTax[1]) allow most people to file their tax returns faster and cheaper by themselves. Similarly, e-commerce websites help cut out wholesalers to buy products and services directly from

* yihan@mis.nsysu.edu.tw

[1] Intuit TurboTax https://turbotax.intuit.com

manufacturers at a lower price. In data visualization, recently a suite of self-service tools emerged that makes it easier for non-technical users to analyze data and generate interactive visualizations themselves. For example, Orange[2] provides flow-based programming that supports building a data analysis and visualization pipeline through combining functional blocks; datapine[3] provides an intuitive drag and drop interface to generate charts; Tableau[4] takes a step further by supporting visualization generation with natural language interfaces (NLI). Recently, IBM Cognos Analytics[5] even support automatic insight discovery and narration on visualizations. As complex decision making can involve a group discussion, collaborative features such as visualization sharing and storytelling, are also taking part in tools such as MS Power BI[6]. These features are for relatively simple analysis and visualization, but it is likely that a few years later when more advanced analytics and visualization features are developed in these systems, the need for professional analysts would begin to drop. As a result, I envision a near future where an intelligent interactive visualization system may take the jobs of many analysts in the decision-making process when the analyses and visualizations are relatively simple.

3 Future Scenario in Data-Driven Decision-Making: Using an Intelligent Interactive Visualization System

Grace, the decision maker of a company, arrives at a meeting room where a group of consultants await. Tom and Ken are sitting at the meeting table and Julie is connected online. Today they need to decide on whether to invest a million dollars in developing a new technology in house to diversify their business. Tom, a junior executive, kicks things off by asking the visualization system, "show us the potential market of (this new technology)." The screen immediately displays a line chart of this technology's global market forecast over the next 10 years. In addition, he also asks the system to provide a financial breakdown of the proposed technology in a parallel coordinates visualization (Figure 2a). Tom examines these visualizations and presents his take on how much of the market he estimates the company can reach in the next few years. As Tom speaks, the system automatically highlights the visuals in the charts to help the group follow his presentation.

All the participants have a tablet or a computer that is connected to the visualization system. Ken, who is a senior executive in the company, is not familiar with the parallel coordinates visualization. But instead of interrupting the presentation or pretending he understands the visualization, he clicks a button next to the parallel coordinates on his tablet to view a quick tutorial on how to read the visualization (Figure 2b). The entire process takes only a few clicks and helps Ken rejoin the meeting without missing much. He then voices his support of Tom's opinion.

On the other hand, Grace questions Tom's estimations of future market share. But instead of using her intuition to challenge Tom, she asks the visualization system to display their company's current market share of a similar technology next to Tom's charts to make her point. The market share of the similar technology in Grace's chart is significantly lower than Tom's estimation (Figure 2c). During this exchange, the system automatically documents these discussions on the sidebar.

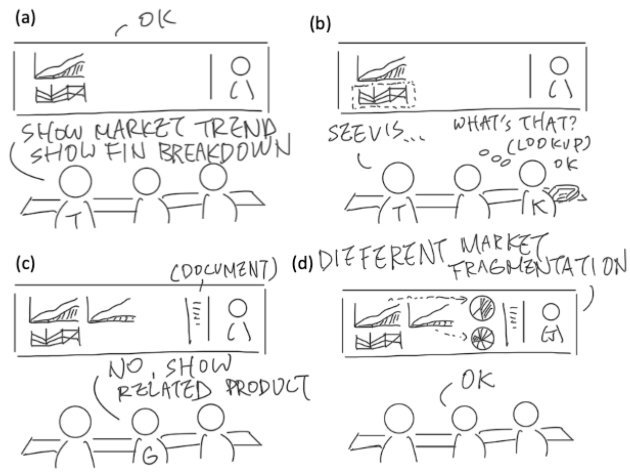

Figure 2: Future Scenario Illustration. (a) Tom asks the visualization system for market trend and financial breakdown of the proposed new technology. (b) Ken uses tablet to learn about unfamiliar visualization. (c) Grace asks for market trend of a similar technology to question Tom's estimation. (d) Julie asks the system for market shares of both technologies to convince Grace that these technologies are different.

Julie, who is an external consultant, steps in with her thoughts. She asks the system to visualize the global market shares of both these technologies in two pie charts. The charts clearly show that the similar technology's market fragmentation is much higher, indicating a tougher competitive landscape. On the other hand, the proposed technology shows less fragmentation (Figure 2d). Therefore, a successful development of the technology has a good chance in capturing a large segment of the market. Grace is convinced by Julie's charts and accepts Tom's market share estimations.

The group continues to discuss the technical aspects and required resources, which again are backed up with data visualizations on the fly in the same way. By the end of the meeting, Grace is convinced that she has enough data to support this million-dollar decision and approves the new technology development.

In this future scenario, I described a data-driven decision-making process without a professional analyst in the loop. All data queries, analyses, and visualizations are made on the fly and could be performed by anyone in the meeting, as shown in Figure 3. Therefore, this type of decision process ensures that the data analysis and visualization are minimally influenced by data analysts.

Nevertheless, the future scenario uses relatively simple analyses and visualizations. It assumes that all the data are well prepared, and the decision team trusts the data and the generated visualizations. For some companies, this decision may even be considered as having relatively low risk.

[2] Orange https://orangedatamining.com
[3] datapine https://www.datapine.com
[4] Tableau https://www.tableau.com
[5] IBM Cognos Analytics https://www.ibm.com/nz-en/products/cognos-analytics
[6] Microsoft Power BI https://powerbi.microsoft.com

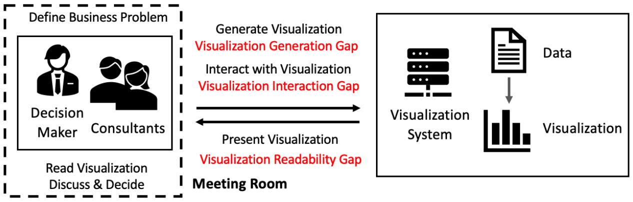

Figure 3: Illustration of the future scenario where all decision-making participants could actively generate and interact with data visualizations. Three gaps need to be addressed before this future is realized.

4 Gaps in Realizing the Future Scenario

However, even to realize such a relatively simple future scenario, several important gaps need to be addressed first. I will focus on the visualization aspect of the system and assume that analysis can be performed automatically in the background. These gaps are the visualization generation gap, the visualization interaction gap, and the visualization readability gap, as shown in Figure 3.

4.1 Visualization Generation Gap

Generating visualizations to this day is not straightforward. Even generating a standard line chart or bar chart in most office productivity tools requires the proper selection and specification of data and manual adjustment of formats. Using advanced business intelligence (BI) tools could help overcome some of these issues, but the complexity of these tools calls for a trained expert. Few tools can generate visualizations fast and easy enough by a non-technical user. As a result, visualizations are often not generated on the fly by the decision team, but are preprepared by well-trained professional analysts who need to guess what the decision team wants to see. This practice may seem convenient to the decision team, but it requires great trust in the data analysts to generate well-designed and unbiased visualizations.

The first potential solution that may fill this gap is to use natural language-based visualization generation systems. However, this technology has a long way to go to be practically usable by this user group. Recent research work provided toolkits (NL4DV) [19], conversational techniques [25], and even deep learning-based systems [6]. Commercial product Tableau has a feature called "Ask Data"[7] that supports natural language queries. However, most of these NLIs require a specific syntax for generating visualizations that is not really "natural." Furthermore, in real-world visualization-based discussions, questions can be in different languages, domain-specific, and vague that require a smart analyst to "decode" and generate the right visualizations. For example, if a decision maker says, "show me how our company is doing recently," the system needs to be smart enough to know which data and visual representations this user is expecting. The term "recently" also requires an inference. If the verbal information is not enough, the NLI might need to provide multiple alternative results or ask for a clarification. These additional inquiries for clarifications by the system need to be clear, succinct, and infrequent to avoid annoying the users. As a result, current tools that require users to spell out the specific data and charts they need are still too tedious and insufficient for practical uses.

The second potential solution that may fill this gap is with task-oriented visualization generation [26]. In this solution, users can select domain-specific tasks from a list to generate visualizations in the given domain. In market research, the list can include "show market share of product X" or "show 5-year market trend of technology Y". From this task list, the decision team will be able to specify what they want more efficiently by recognizing the tasks they prefer than recalling them from memory [20]. The list can be sorted by frequency of use to make it easier to find commonly performed tasks. Selecting a task can generate relevant visualizations from a pretrained knowledge base that ensures the visualizations are appropriate for the task. For example, market shares of a technology should be visualized as pie charts because it is the accepted practice. The challenge is that these domain-specific task-visualization mappings need to be predefined and maintained by current data analysts. An alternative way to define these mappings is to acquire them from other decision teams. But it is unknown whether mappings from other teams are suitable for the current team and whether other decision teams are willing to share this knowledge.

4.2 Visualization Interaction Gap

Interactive features that allow users to manipulate visualization views on demand are invaluable to support the dynamic decision-making needs. However, they are underutilized. From Schneiderman's information visualization mantra, "overview first, zoom and filter, then details on demand," it is clear that interactions are at the heart of an effective visualization system [28]. Over the years, many novel visualization interaction techniques were developed [36]. These interactions were incorporated into interactive visualization tools such as Tableau and MS Power BI. But currently most visualizations or charts created and used for decision making are static, not interactive [7]. This practice significantly limits their potential value for data-driven decisions.

I believe the lack of use of interactive visualizations in decision making is not because interactions are not useful but because they are not easy enough to use for decision teams. For example, zooming and panning of visualizations on a desktop require the familiarity of combining control+scrolling and dragging. None of these techniques are considered "visible" to users if we consider Norman's design principles [21]. Some visualizations have a fancy control panel that supports powerful filtering, formatting, and data manipulation capabilities. However, without a significant amount of training and practice, I doubt any user will be comfortable performing these interactions live in front of a crowd, especially when this crowd is full of superiors who are unusually stressed and impatient.

A potential solution to this problem is to make visualization interactions respond to an NLI as described in the future scenario. A user could verbally indicate what to highlight, what to filter, and where to zoom in to a system and the system would perform these visualization interaction tasks for the user, as if a human analyst is in the room. Decision makers are used to giving verbal commands to communicate their needs. Therefore, having a system that can understand these verbal commands provides a familiar interface to this user group. This NLI should be easier to develop than the visualization generation interface as the potential commands for interactions are constrained by the presented visualization. But a verbal interface for interacting with visualizations has downsides as well. For example, it may be difficult to verbally indicate a specific circle to select among a large group of circles in a scatterplot.

To overcome this type of selection problem in an NLI is to provide other modes of interactions that allow decision teams to interact with visualizations more directly. For example, if they can

[7] Tableau Ask Data: https://help.tableau.com/current/pro/desktop/en-us/ask_data.htm

use their fingers to point at what they want to select on the screen, it will be much easier than using words to do so. The challenge is to have the right technology to identify the finger in midair and map it to the location on screen. Alternatively, if the visualization is mirrored on remotely-connected tablets, the decision team could interact with the visualization on the tablets to accomplish the same task. The downside is that decision team members will have to lower their heads to operate the tablets that may interrupt the interaction with other meeting participants.

4.3 Visualization Readability Gap

Even when visualizations can be easily generated and manipulated for presentation needs, they are still not valuable if their intended users cannot read them. Researchers have studied the visualization readability issue for many visual representations [1, 3, 24, 29, 32] and developed tests for assessing visualization literacy [15]. Based on the literature, it seems that visualizations are not readable for three reasons. First, the visual representation is inherently complex, such as a graph with a large number of nodes and links that look like a giant fur ball [29]. Second, some individuals may have less innate ability in reading visualizations, such as those with poor spatial abilities [32]. Third, most people are not sufficiently trained to read a wide variety of charts and graphs beyond statistical graphs learnt in our formal educations [33]. With the increase in variety and volume of data, advanced visualizations are needed to effectively display such multivariate data. However, reading these visualizations poses an even greater challenge to non-technical decision teams [22].

For example, parallel coordinates can effectively visualize multidimensional data at scale but is notoriously difficult to read [11]. As a result, they are not typically used in corporate decision-making environments. However, this lack of use is why I selected it in the future scenario. This visualization technique is proven to be very effective with 5000+ publications in the research literature surveyed about a decade ago [11]. But they are usually not taught in business schools and require advanced interactions (filtering & reordering axes) to effectively explore. With the availability of big data, emerging business tasks may need advanced visualizations to support [22]. Hence, I believe that in the future, this type of advanced multivariate visualizations could find their place in data-driven decision-making meetings.

To fill the readability gap, many researchers redesigned or simplified visualizations. For example, node-link graphs can be redrawn with less edge crossings or motifs to improve aesthetics and readability [9, 23, 24, 29] and temporal events can be simplified through filtering and aggregation [18]. Although these methods are useful for improving the readability of graphs and temporal visualizations, they are not easily applicable to other types of visualizations, such as parallel coordinates. Therefore, to help decision teams who need to read a wide variety of visualizations, there should be a more generalizable method that helps improve the readability of these visualizations. Ideally, this help can be provided in a form of training that can be easily conducted on the fly the first time the visualizations are seen.

In the future scenario, a quick training is provided on a tablet to the user in need. This type of training can also be provided on stage to make sure everyone can understand the presented visualization. Some potential designs for this type of training can be found on websites where interactive charts are explained with animations of carefully crafted tutorials (e.g., Gapminder's How to Use [8]). However, these designs' effectiveness and generalizability have not been carefully validated. Therefore, how to provide an effective, on-the-fly training method to improve the readability of a variety of visualizations is still an open question.

5 Discussion

Although there is much potential in using intelligent interactive visualization systems for decision making, a few questions arise on whether the future scenario presented in this paper could be realized.

5.1 Will decision teams generate unsuitable visualizations that result in poor analysis outcomes?

Professional analysts today do significantly more in data processing, analysis, and visualization than what was presented in the relatively simple future scenario. Their expertise ensures that a good selection of analysis and visualization methods are performed to generate a quality analysis outcome. However, I believe it is possible to infuse much of this knowledge into visualization systems so that the decision teams can generate acceptable data visualizations and reasonable analysis results by themselves in many practical cases. This knowledge is domain specific because, for example, charts that are acceptable and commonly used in accounting decisions are very different from those used in manufacturing decisions [22, 37]. Even within the same application domain, different decisions based on different data may also need different visualizations. These acceptable analysis and visualization knowledge should be taught to the intelligent visualization systems by current data analysts.

However, how do data analysts effectively transfer their knowledge of data analysis and visualization into a system? There is only a limited number of analysts in a given company, how could they generate enough training data for the system to learn? An idea is to take lessons from the AI and recommendation system communities. In AI, there are pretrained models that can be modified or "transferred" for specific uses [34]. I can imagine someone pretraining a set of visualization generation models for others to transfer into more domain-specific visualization generation models. Visualization generation can also be modelled as a visualization recommendation task. In recommendation systems, there is a well-known technique called collaborative filtering that could recommend items that are preferred by similar users. A visualization system could also have such a function build-in that can identify similar decision team's data analysis and visualization practices and recommend them to the current decision team. In these cases, other decision teams' practices will have to be shared first and modelled into the system. Nevertheless, some companies may consider these analysis practices as their trade secret. As a result, how to share this information without compromising confidentiality will be a challenge.

In addition, the visualization system should provide several intelligent operations to ensure that important insights are not overlooked. For example, when a decision team specifies a task, the system should automatically recommend not one, but multiple visualizations that could provide different views of the data. To help decision teams capture the various insights in every visualization, AI could help automatically point them out through annotations and labels [12]. Insights could be identified by visual patterns such as the scagnostic measures (e.g., outliers, shape) developed for evaluating scatterplot matrices [35]. These automatically identified insights together could form a less-biased, data-driven story given by an AI system than by a human storyteller.

[8] Gapminder: https://www.gapminder.org/tools

5.2 Will we ever have an NLI that can interact with the decision team just like a human analyst?

In the future scenario, I described many interactions with NLI. NLI has come a long way in understanding human speech and carrying out conversations [27]. A realistic NLI system should be one that passes the "Turing test" where the users won't realize if they are talking to a computer [31] Currently, NLI is good enough for transcription tasks, but guessing users' intents and following a flow of conversation like a human analyst is still very difficult. This is because a real person has so much more knowledge, context, and experiences in understanding one another, verbally and non-verbally. Therefore, a good NLI needs to understand not only the verbal information, but also the who, what, where, when of the asked question to correctly understand the context and generate the right visualizations. Technology-wise, this goal might be achievable in the future with ubiquitous sensing in place. But would we ever be comfortable providing this much information to a computer about ourselves? If it generates too much privacy concerns, will such an NLI system ever have enough information to be successful?

5.3 Will data collected ever have sufficient quality to take such a big role in decision making?

Data are messy, scattered, and difficult to capture. Many valuable visualizations require high-quality, digitized, and sufficient data. To have this type of data, it will require significant data preparation, processing, organization, and maintenance efforts, which will need an army of data collectors and data engineers. Therefore, this process obviously needs to be automated to streamline and scale such operation.

Nevertheless, the more data are collected, the more effort it takes to process the data to ensure it remains of high quality. Will a company be willing to invest in such a great endeavor? For many private companies, data cannot be shared outside the organization. Therefore, will collecting such data only for their own use ever be cost efficient?

5.4 Will non-technical decision teams ever be comfortable with operating visualization systems and trusting the visualized results?

Only having the tools ready may not be enough to ensure that decision teams can self-serve their analytic and visualization needs. Decision team members often come from various backgrounds that may be reluctant to operate novel technical systems themselves. The good news is that the new generation of decision teams are more familiar with operating computing systems than the previous ones [13]. Furthermore, more professionals than ever are taking advantages of the global availability of online courses and certificates to learn about data analytics and visualizations. For example, as of August 2022, one of the top data analysis courses on Coursera (Google Data Analytics Professional Certificate) that include training on both data analysis and visualizations, has enrolled over one million students [10]. It is not hard to imagine that with this many people being trained, in a few years we may have many decision teams that are well-prepared to operate visualization systems for their own data-driven decision needs.

On a related concern, decision teams may distrust the visualized results generated by the visualization system. The level of distrust may depend on the decision teams' familiarity with the data and the visualization system. However, as decision teams are being better trained in analytics and visualizations, the amount of distrust could be reduced over time. In addition, this concern could be further mitigated when we are giving decision teams domain-specific visualizations with which they are already familiar and providing them interactive control to easily configure the generated results.

5.5 If this type of future scenario comes true at a large scale, what would data analysts likely do in the future?

Data analysts may slowly transition into different roles as in any technology-driven paradigm shifts. They could first be excused from meetings for simple, low-risk data-driven decisions, but later move out of more complex, high-risk decision-making processes. Specifically, many of today's data analysts may be transferred into working in the background with data engineers to direct how data should be prepared and managed as well as how interactive visualization systems should be designed to support decision teams. For example, in task-oriented visualization generation, data analysts will need to transfer their expertise and knowledge on what decision teams may want to see for any task into the visualization system. These new roles are not easy and require the thorough understanding of the current data-driven decision-making process. As different organizations have different processes and data, they will need different visualization system designs to generate the best views of the data for their decisions.

However, a small set of analysts may still help with the analysis and visualization of decisions that call for uniquely complex analyses. The need for analysts in these cases may be because these analyses may require so much expertise that they could never be cost-efficiently incorporated into a visualization system.

6 Conclusion

In this paper, I envision a future scenario where professional analysts may not be needed for data analysis and visualization in the data-driven decision-making process. Decision teams will be able operate intelligent interactive visualization systems on their own to acquire the data they need. However, even when the analyses and visualizations are relatively simple, there are three gaps: the visualization generation gap, the visualization interaction gap, and the visualization readability gap, that need to be addressed to realize this future scenario. Furthermore, I proposed potential solutions to these gaps and discussed many other related questions. I hope this paper can raise an interest and start a discussion in the community on how future intelligent interactive visualization systems could be designed to best support non-technical decision teams in making data-driven decisions.

7 Acknowledgements

I would like to thank the reviewers for their thoughtful comments and suggestions. This work was supported in part by the National Science and Technology Council in Taiwan (MOST 111-2410-H-110-004-MY2) and the Intelligent Electronic Commerce Research Center from The Featured Areas Research Center Program within the framework of the Higher Education Sprout Project by the Ministry of Education (MOE) in Taiwan.

References

[1] Bennett, C. et al. 2007. The Aesthetics of Graph Visualization. *Computational Aesthetics in Graphics, Visualization, and Imaging* (2007), 57–64.

[2] Berntsson Svensson, R. and Taghavianfar, M. 2020. Toward Becoming a Data-Driven Organization: Challenges and Benefits. *Lecture Notes in Business Information Processing*. 385 LNBIP, (Sep. 2020), 3–19. DOI:https://doi.org/10.1007/978-3-030-50316-1_1.

[3] Börner, K. et al. 2015. Investigating aspects of data visualization literacy using 20 information visualizations and 273 science museum visitors: *Information Visualization*. 15, 3 (Aug. 2015), 198–213. DOI:https://doi.org/10.1177/1473871615594652.

[4] Chuang, J. et al. 2012. Interpretation and trust: Designing model-driven visualizations for text analysis. *Conference on Human Factors in Computing Systems - Proceedings*. (2012), 443–452. DOI:https://doi.org/10.1145/2207676.2207738.

[5] Dasgupta, A. et al. 2016. Familiarity vs trust: A comparative study of domain scientists' trust in visual analytics and conventional analysis methods. *IEEE transactions on visualization and computer graphics*. 23, 1 (2016), 271–280.

[6] Dibia, V. and Demiralp, Ç. 2019. Data2vis: Automatic generation of data visualizations using sequence-to-sequence recurrent neural networks. *IEEE computer graphics and applications*. 39, 5 (2019), 33–46.

[7] Dimara, E. et al. 2021. The Unmet Data Visualization Needs of Decision Makers within Organizations. *IEEE Transactions on Visualization and Computer Graphics*. (2021). DOI:https://doi.org/10.1109/TVCG.2021.3074023.

[8] Dörk, M. et al. 2013. Critical InfoVis: exploring the politics of visualization. *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 2189–2198.

[9] Dunne, C. and Shneiderman, B. 2013. Motif simplification: Improving network visualization readability with fan, connector, and clique glyphs. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), 3247–3256.

[10] Google Data Analytics Professional Certificate | Coursera: *https://www.coursera.org/professional-certificates/google-data-analytics*. Accessed: 2022-08-30.

[11] Heinrich, J. and Weiskopf, D. 2013. State of the Art of Parallel Coordinates. *Eurographics 2013 - State of the Art Reports* (2013).

[12] Hullman, J. et al. 2013. Contextifier: automatic generation of annotated stock visualizations. *Proceedings of the SIGCHI Conference on human factors in computing systems* (2013), 2707–2716.

[13] IDG 2016. *Data & Analytics Survey*.

[14] Kandogan, E. et al. 2014. From data to insight: work practices of analysts in the enterprise. *IEEE computer graphics and applications*. 34, 5 (2014), 42–50.

[15] Lee, S. et al. 2017. VLAT: Development of a Visualization Literacy Assessment Test. *IEEE Transactions on Visualization and Computer Graphics*. 23, 1 (Jan. 2017), 551–560. DOI:https://doi.org/10.1109/TVCG.2016.2598920.

[16] Madanagopal, K. et al. 2019. Analytic Provenance in Practice: The Role of Provenance in Real-World Visualization and Data Analysis Environments. *IEEE Computer Graphics and Applications*. 39, 6 (2019), 30–45. DOI:https://doi.org/10.1109/MCG.2019.2933419.

[17] Mayr, E. et al. 2019. Trust in Information Visualization. *EuroVis Workshop on Trustworthy Visualization (TrustVis)* (2019).

[18] Monroe, M. et al. 2013. Temporal event sequence simplification. *IEEE Transactions on Visualization and Computer Graphics*. 19, 12 (2013), 2227–2236. DOI:https://doi.org/10.1109/TVCG.2013.200.

[19] Narechania, A. et al. 2020. NL4DV: A toolkit for generating analytic specifications for data visualization from natural language queries. *IEEE Transactions on Visualization and Computer Graphics*. 27, 2 (2020), 369–379.

[20] Nielsen, J. 1994. Enhancing the explanatory power of usability heuristics. *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (1994), 152–158.

[21] Norman, D. 2002. *The design of everyday things*. Basic books.

[22] Perkhofer, L.M. et al. 2019. Interactive visualization of big data in the field of accounting: A survey of current practice and potential barriers for adoption. *Journal of Applied Accounting Research*. (2019).

[23] Purchase, H. 1997. Which aesthetic has the greatest effect on human understanding? *Graph Drawing* (Berlin, Heidelberg, 1997), 248–261.

[24] Purchase, H.C. 2002. Metrics for Graph Drawing Aesthetics. *Journal of Visual Languages & Computing*. 13, 5 (Oct. 2002), 501–516. DOI:https://doi.org/10.1006/JVLC.2002.0232.

[25] Setlur, V. et al. 2016. Eviza: A natural language interface for visual analysis. *Proceedings of the 29th annual symposium on user interface software and technology* (2016), 365–377.

[26] Shen, L. et al. 2021. TaskVis: Task-oriented Visualization Recommendation. *Proc. EuroVis* (2021).

[27] Shen, L. et al. 2021. Towards Natural Language Interfaces for Data Visualization: A Survey. *arXiv preprint arXiv:2109.03506*. (2021).

[28] Shneiderman, B. 2003. The eyes have it: A task by data type taxonomy for information visualizations. *The craft of information visualization*. Elsevier. 364–371.

[29] Sugiyama, K. 2002. *Graph Drawing and Applications for Software and Knowledge Engineers*. World Scientific.

[30] The Advantages of Data-Driven Decision-Making: 2021. *https://online.hbs.edu/blog/post/data-driven-decision-making*. Accessed: 2022-01-27.

[31] Turing, A.M. 2009. Computing machinery and intelligence. *Parsing the turing test*. Springer. 23–65.

[32] Velez-Rojas, M.C. 2009. *A cognition-based framework for the development of visualization literacy*. Rutgers The State University of New Jersey-New Brunswick.

[33] Wallgren, A. et al. 1996. *Graphing statistics & data: Creating better charts*. Sage.

[34] Weiss, K. et al. 2016. A survey of transfer learning. *Journal of Big data*. 3, 1 (2016), 1–40.

[35] Wilkinson, L. et al. 2005. Graph-theoretic scagnostics. *Information Visualization, IEEE Symposium on* (2005), 21.

[36] Yi, J.S. et al. 2007. Toward a deeper understanding of the role of interaction in information visualization. *IEEE transactions on visualization and computer graphics*. 13, 6 (2007), 1224–1231.

[37] Yusof, E.M.M. and Othman, M.S. 2012. A Review on the Dashboard Characteristics for manufacturing Organizations. *Journal of information systems research and innovation*. 2, (2012), 28–34.