

TOWARDS A FRAMEWORK FOR AI-ASSISTED DATA STORYTELLING

Angelica Lo Duca¹

*Institute of Informatics and Telematics, National Research Council, via G. Moruzzi 1, 56124 Pisa, Italy
angelica.loduca@iit.cnr.it*

Keywords: Data Storytelling, Generative AI, Data Visualization

Abstract: Data storytelling is building stories supported by data to engage the audience and inspire them to make decisions. Applying data storytelling to data visualization means adding a narrative that better explains the visual and engages the audience. Generative AI can help transform data visuals into data stories. This paper proposes AI-DaSt (AI-based Data Storytelling), a framework that helps build data stories based on generative AI. The framework focuses on visual charts and incorporates two main generative AI models provided by the OpenAI APIs: text generation and image generation. We use GPT-3.5 for the chart title, commentary and notes, and image generation for images to include in the chart. We also describe the potential ethical issues and possible countermeasures related to using Generative AI in data storytelling. Finally, we focus on a practical use case, which shows how to transform a data visualization chart into a data story using the implemented framework.

1 INTRODUCTION

Data storytelling builds stories supported by data, allowing analysts and data scientists to present and share their insights to engage the audience and inspire them to make decisions. Data storytelling is used for different purposes, such as business (Knafllic 2015) and education (Ma 2012). A visual data story (data story, for short) combines graphs, words, and images (Kosara 2013) in a sequence of elements: beginning, middle, and end (Bordwell 2003).

A data story comprises three main aspects: data, visuals, and narrative (Dykes 2019). Data is the building block of each data story, visual is how data is represented, such as graphs, infographics, videos, etc. The narrative is the story built around data to guide the audience to understand the meaning of what's being shown.

In recent years, generative AI (Gozalo-Brizuela 2023) has opened up new possibilities for enhancing data visualization with narrative elements. Generative AI is a subfield of Artificial Intelligence (AI), that uses Large Language Models (LLM) to generate new text based on a given input, called prompt.

This paper introduces AI-DaSt (AI-based Data Storytelling), a novel framework that leverages generative AI tools to add a narrative to a data

visualization chart. Specifically, our framework incorporates LLM and generative image models provided by the OpenAI API. We use GPT-3.5 (Brown 2020), a generative text model, to generate chart titles, commentary, and annotations and DALL-E (Ramesh 2021) to generate visually appealing images to be added to the data visualization chart.

To illustrate the practical application of our framework, we present a detailed use case. This use case highlights how AI-DaSt empowers analysts and data scientists to transform a data visualization chart into a data story. Our proposed AI-DaSt framework could empower data visualization charts with engaging narratives.

The paper also describes the potential ethical issues and countermeasures related to the usage of generative AI in data storytelling, focusing on potential misinformation and bias.

The remainder of the paper is organized as follows. In Section 2, we describe the related literature focusing on data visualization, data storytelling, and generative AI. Section 3 describes our methodology to convert a data visualization chart into a data story. Section 4 focuses on the potential ethical issues and countermeasures related to generative AI in data storytelling. Next, Section 5 illustrates the AI-DaSt user interface, and Section 6

¹ <https://orcid.org/0000-0002-5252-6966>

the use case. Finally, in Section 7, we give our conclusions and future work.

2 RELATED WORK

The literature associated with this work relates to three main research fields: data visualization, data storytelling, and generative AI.

2.1 Data Visualization

The literature about data visualization is rich in principles, techniques, and tools for building appealing data visualization charts (Qin 2020). Research in data visualization has focused on the design aspects (Borner 2015), cognition process (Ware 2013), and technical aspects (Wilkinson 1999), meaning that visual styles of data storytelling are often designed simplistically (Lee 2015).

A vast literature exists on automatically generating texts summarizing statistics and trends described by a visual chart. We can classify these works into two categories: those using Computer Vision (CV) and Natural Language Process (NLP) (Chen 2020, Kantharaj 2022, Mittal 1998, Qian 2021, Srinivasan 2018), and those using structured specifications of the chart’s construction.

CV-based approaches build a CV model that learns the salient features of rasterized images taken as an input and produces a relevant text as an output (Team 2014, Brown 2011, Satyanarayan 2016). However, even the most mature tools, such as Tableau Public (Milligan 2022) and Microsoft Power BI (Ferrari 2016), do not provide any feature to build the actual narrative depicted by a chart.

All the described techniques can generate captions for charts, but they cannot build engaging text that transforms them into an engaging story for the audience. Compared to the current literature, this paper aims to define a first tentative to generate engaging texts and images to incorporate into a chart using generative AI.

2.2 Data Storytelling

The literature about data storytelling is varied and covers different aspects, such as the role of rhetoric in building narratives (Hullman 2011, Hullman 2013). A great effort has been made to identify commonly used approaches to build stories in the media and news field. Segel and Heer propose seven genres of narrative visualization for newspaper stories: magazine style, annotated chart, partitioned poster, flow chart, comic strip, slide show, and video (Segel 2010). The approaches proposed by Segel and Heer

are limited to the specific scenario of newspapers. Other scenarios would require additional approaches and critical evaluation of the effectiveness of the built stories (Kosara 2013).

Lundgard & Satyanarayan organized the semantic content of textual descriptions of charts into four levels: enumerating visualization construction properties, reporting statistical concepts and relations, identifying perceptual and cognitive phenomena, and explaining domain-specific insights (Lundgard 2021).

Compared to the existing literature, we propose a tool that assists users in building chart annotations using generative AI.

2.3 Generative AI

Although generative AI is a very young research field, many works exist in the literature, covering a variety of tasks from storytelling (Akoury 2020, Nichols 2020) to code synthesis (Austin 2021) and email auto-completion (Yonghui 2018).

AI for storytelling is used mainly in education (Ali 2021, Crompton 2022, Han 2023) and co-writing (Yuan 2022).

Compared to the current literature, we focus on how to add a narrative to a data visualization chart in terms of annotations. We introduce the use of generative AI to enrich data visualization charts.

3 FROM DATA VISUALIZATION TO DATA STORYTELLING

Data storytelling tools can be classified into different categories based on the type of feature provided: annotated chart, timeline and storyline, data video, scrolly-telling, data comics, and map (Ren 2023). This paper focuses on annotated charts and how to transform a raw chart into a data story through engaging annotations. We suppose we already have a basic chart representing an insight extracted from data, and we want to add annotations that help the audience understand the chart. Annotations may vary based on the audience reading the chart (Lundgard 2021).

In the remainder of this section, we will describe the types of annotations, the different types of audiences, and the generative AI tools used to help build annotations adapted to the audience type.

3.1 Annotation Types

The English Oxford Dictionary defines an annotation as *a note by way of explanation or*

comment added to a text, document, diagram, etc., or to a particular copy of a text, document, diagram, etc.². In this paper, we define an annotation as a text or an image that helps the reader set the context behind data and understand the chart's meaning. In their work, Stokes *et al.* argue that people prefer charts that also include text and, in particular, charts with a great presence of annotations (Stokes 2022).

We consider four types of annotations:

- *Title* - a concise phrase that summarizes the primary purpose of the chart. How a title is written influences how a chart is interpreted (Kong 2018);
- *Commentary* - a brief comment that provides additional information or context to the chart;
- *Note* - a brief written remark providing additional information on a specific point in the chart.
- *Image* - a visual representation of the main subject of the chart. We can use images to engage the audience from an emotional perspective.

3.2 Audience

The audience is the person or group reading a chart. Understanding the target audience is crucial to building data stories that convey information effectively (Dykes 2019). This paper focuses on three types of audiences: the general public, executives, and professionals.

The general public comprises individuals from various backgrounds and levels of knowledge. They may have little to no previous knowledge of the chart subject. When crafting data stories for the general public, we should use precise language, avoid overwhelming them with too much information, and focus on presenting the most relevant insights visually and engagingly.

Executives are typically high-level decision-makers in organizations who rely on data-driven insights to make essential business choices. They often have limited time and need concise and actionable information. When creating data stories for executives, it is essential to present key findings, trends, and recommendations upfront. We should use visualizations highlighting the most critical data points and providing a straightforward narrative linking the data to strategic goals.

Professionals consist of individuals with a specific domain expertise or professional background. They have a deeper understanding of data and require more analytical information. When creating data stories for professionals, we should explain the data analysis's methodology,

assumptions, and limitations. We should also consider including additional supporting data and references, allowing professionals to explore the data further.

3.3 Generative AI Models

Generative AI models can be used for different purposes, including text generation (chatbots, creative writing, and code generation) and image and audio generation.

This paper focuses on GPT-3.5, an advanced version of the Generative Pre-trained Transformer 3 (GPT-3) model. GPT3.5 generates text in tasks like chatbot interactions, creative writing, and code generation.

To generate images, we use DALL-E, a generative AI model developed by OpenAI that combines the power of (Generative Adversarial Network) GANs and transformers to generate images from textual descriptions.

4 AI-DaSt

AI Data Storytelling, AI-DaSt for short, is a web framework enabling users to improve their data visualization charts through generative AI. Figure 1 shows the AI-DaSt architecture, which comprises three main elements: the Main Editor, the Textual Annotator, and the Graphic Annotator.

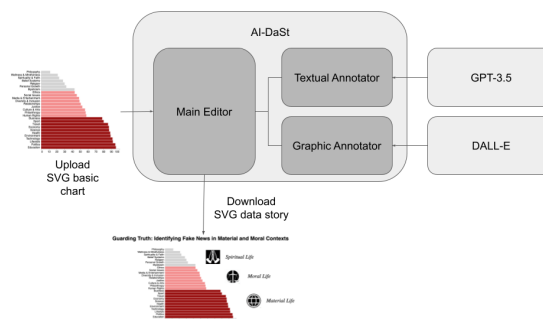


Figure 1 The AI-DaSt architecture

4.1 The Main Editor

The Main Editor receives as input a Scalable Vector Graphics (SVG) image, which already contains a basic chart, and produces another SVG image as an output. Different tools can be used to generate the SVG basic chart, such as the Altair Python library

² <https://www.oed.com/search/dictionary/?scope=Entries&q=annotation>

(VanderPlas 2018). The Main Editor provides the following functionalities:

- Add/modify the chart title
- Add/modify the chart commentary/note
- Add graphics to the chart
- Export the chart as an SVG image.

To implement the first two functionalities, the Main Editor uses the Textual Annotator, and the third functionality, uses the Graphic Annotator. To activate a functionality, the Main Editor provides a button that opens a popup window.

4.2 The Textual Annotator

The Textual Annotator provides an interface to the OpenAI API to automatically generate texts that are added to the SVG chart. We use the GPT-3.5 model to generate texts. When connecting to the OpenAI API, the Textual Annotator sets the following parameters (in addition to the OpenAI key):

- The *text type*: one among title, description, and annotation
- The *audience type*: one among the technical audience, general audience, and executives
- *Number of produced outputs*: a number between 1 and 5
- The *maximum number of characters to generate*: a number between 8 and 250
- The *description*: the topic of the chart. The model will generate the output based on this field.

The following piece of pseudo-code defines the basic structure of a prompt used to instruct GPT-3.5:

```
Generate <Number of outputs> <text types> for
<audience type> on the following topic, using
max <N> characters: <description>
```

For example, we can write the following prompt:

```
Generate 3 titles for technical experts on
the following topic, using max 200
characters: the product sales over the last
10 months have increased of 72%.
```

Figure 2 shows the Textual Annotator mockup.

4.3 The Graphic Annotator

The Graphic Annotator provides an interface to the DALL-E model provided by the OpenAI API to generate images that are added to the SVG chart automatically. When connecting to the OpenAI API, the Graphic Annotator sets the following parameters (in addition to the OpenAI key):

- *Style* - the artistic representation of the image. It is one among cartoons, illustrations, photography, and icons;
- *Details* - the type of image. It is one among black and white, realistic, ambient light
- *Number of images* to generate;
- *Description* - the topic of the image.

The following piece of pseudo-code defines the basic structure of a prompt used to instruct DALL-E:

```
A <details> <style> about <description>
```

For example, we can write the following prompt:

```
A <black and white> <illustration> about <an
old castle>
```

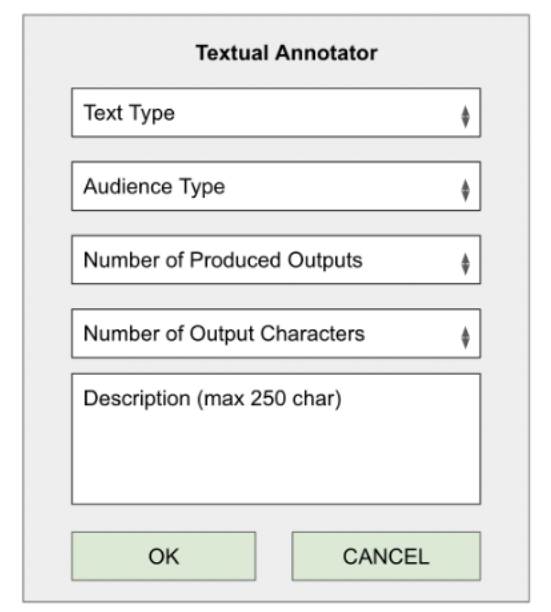


Figure 2 The Textual Annotator mockup

Figure 3 shows the Graphic Annotator mockup.

5 ETHICAL ISSUES

The use of generative AI can raise some ethical issues (Stahl 2024). When using generative AI in data storytelling, we should consider at least two primary ethical issues: bias and misinformation.

Firstly, bias in AI refers to systemic and unjustified preferences, stereotypes, or prejudices in AI systems due to altered training data (Roselli 2019). This can result in narratives that inadvertently perpetuate stereotypes or unfair representations of certain groups, undermining the objectivity and fairness of data stories.

Secondly, misinformation can arise when AI systems generate plausible-sounding content, which does not have any correspondence with reality. This may lead to disseminating misleading information and building totally fake data stories that seem plausible.

To overcome these ethical issues, one approach could be to always review the content produced by generative AI tools and align it with the UNESCO ethical guidelines (UNESCO 2021).

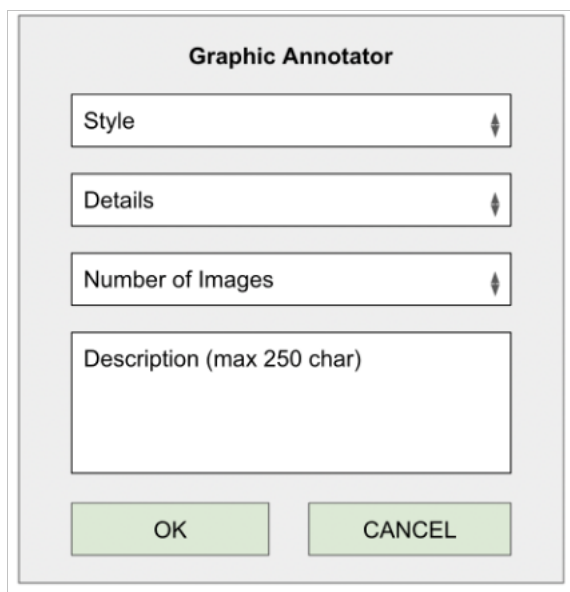


Figure 3 The Graphic Annotator mockup

6 CASE STUDY

As a case study, we consider the following scenario. Let us imagine that XX is an important website that publishes news from different contributors. At a given point, the editor-in-chief receives some complaints from different readers because they read many fake news. The editor-in-chief wants to analyze the number of fake news on the XX website and advise the website editors to pay attention to the categories of news with the highest probability of being fake. The editor-in-chief has already collected data, as shown in Table 1.

Table 1: An extract of the dataset of the example

Category	Number of Fake Articles	Number of Articles
Politics	1235	1300

Economy	1456	1678
Justice	300	570
Religion	30	100

Figure 4 shows a first representation of the dataset.

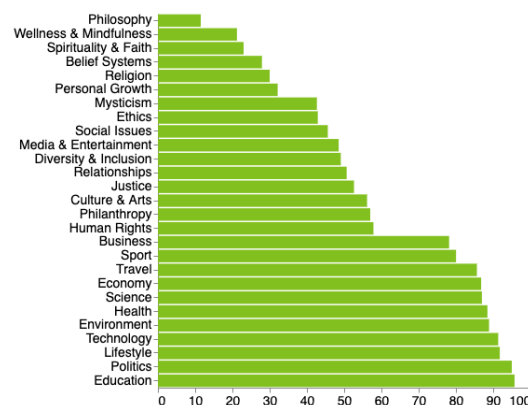


Figure 4 A preliminary chart representation

If we look at labels carefully, we can notice that at the bottom of the pyramid, there are categories related to material life (from Education to Business). In the middle of the pyramid are moral life categories (from Human Rights to Ethics). At the top of the pyramid are categories related to spiritual life (from Mysticism to Philosophy). This means that most fake news relates to material life (more than 70%) and moral life (more than 30% of fake news, but less than 60%). We can highlight the model of material-moral-spiritual life using different colors in the chart based on the different macro categories the news belongs to. Figure 5 shows the resulting chart. We have used two tonalities of red to highlight the urgency of paying attention to material and moral life. We have done this preliminary step manually, without the usage of AI-DaSt.

The next step to transform the chart into a data story involves adding annotations. We give the chart shown in Figure 5 as input to AI-DaSt.

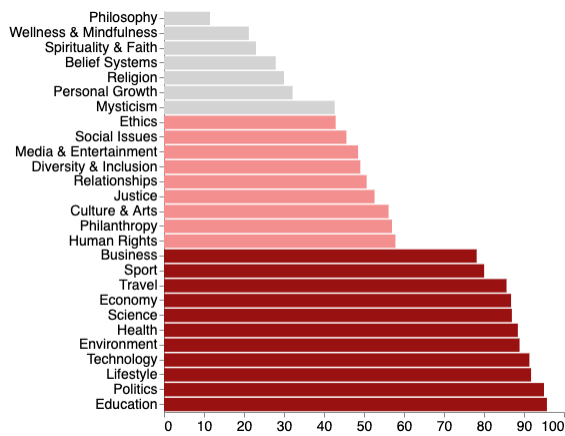


Figure 5 The chart of Figure 4 with macro categories highlighted

First, we add three images, one for each category. We use the Graphic Annotator. For each image, we set the style to icon, the details to black and white, the number of images to four, and the description to praying hands (spiritual life), balance (moral life), and a circle representing the world (material life), respectively. Figure 6 shows the generated images for spiritual life.



Figure 6 The images generated for the spiritual life

Near the icon, we can add a text describing the macro category. Figure 7 shows the resulting chart.

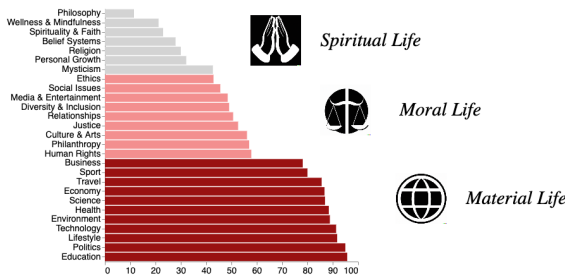


Figure 7 The chart after using the Graphic Annotator

The icons and the text act as the legend. In addition, they are the characters of our story.

The next step involves adding a title to the chart. We set the following parameters:

- The text type: title
- The audience type: technical audience
- Number of produced outputs: 3

- The maximum number of characters to generate: 200
- The description: pay attention to material and moral news because they have a high percentage of fake news.

The Textual Annotator produces the following titles:

1. Navigating the Deceptive Landscape: Analyzing Material and Moral News for Fake Content Detection
2. Guarding Truth: Identifying Fake News in Material and Moral Contexts
3. Cracking the Code of Misinformation: Material and Moral News Authentication Strategies.

We choose the second, and we add it to the chart, as shown in Figure 8.

Guarding Truth: Identifying Fake News in Material and Moral Contexts

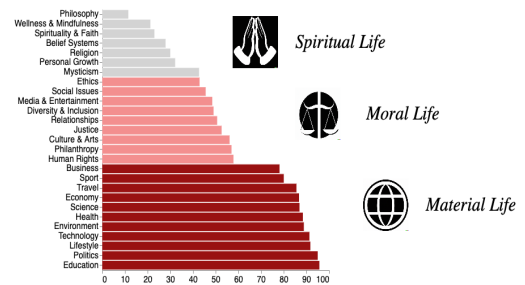


Figure 8 The images generated for the spiritual life

7 CONCLUSIONS AND FUTURE WORK

In this study, we have described AI-DaSt, a framework for turning data visualization charts into data stories through Generative AI. The framework uses GPT-3.5 and DALL-E to generate text and images to include in the chart. At the moment, AI-DaSt is at the design level, and we will implement it in Python.

The framework can potentially empower analysts and data scientists to create engaging data stories from simple charts. Ethical issues should also be considered while building data stories using generative AI. Future work could involve further refining the AI-DaSt framework, addressing scalability challenges, and exploring additional applications of generative AI in data storytelling.

REFERENCES

- Akoury, N., Wang, S., Whiting, J., Hood, S., Peng, N., & Iyyer, M. (2020). Storium: A dataset and evaluation platform for machine-in-the-loop story generation. arXiv preprint arXiv:2010.01717.
- Ali, S., DiPaola, D., Lee, I., Hong, J., & Breazeal, C. (2021, May). Exploring generative models with middle school students. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1-13).
- Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski, H., Dohan, D., ... & Sutton, C. (2021). Program synthesis with large language models. arXiv preprint arXiv:2108.07732.
- Bordwell, D. and Thompson, K. (2003). *Film Art: An Introduction*. McGraw-Hill,
- Borner, K. (2015). *Atlas of knowledge: Anyone can map*. The MIT Press: Cambridge, MA, USA.
- Brown, A., & Wilson, G. (2011). *The Architecture of Open Source Applications: Elegance, Evolution, and a Few Fearless Hacks* (Vol. 1). Lulu. com.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877-1901.
- Chen, C., Zhang, R., Koh, E., Kim, S., Cohen, S., & Rossi, R. (2020). Figure captioning with relation maps for reasoning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (pp. 1537-1545).
- Crompton, H., Jones, M. V., & Burke, D. (2022). Affordances and challenges of artificial intelligence in K-12 education: A systematic review. *Journal of Research on Technology in Education*, 1-21.
- Dykes, B. (2019). *Effective data storytelling: how to drive change with data, narrative and visuals*. John Wiley & Sons.
- Ferrari, A., & Russo, M. (2016). *Introducing Microsoft Power BI*. Microsoft Press.
- Gozalo-Brizuela, R., & Garrido-Merchan, E. C. (2023). ChatGPT is not all you need. A State of the Art Review of large Generative AI models. arXiv preprint arXiv:2301.04655.
- Han, A., & Cai, Z. (2023, June). Design implications of generative AI systems for visual storytelling for young learners. In *Proceedings of the 22nd Annual ACM Interaction Design and Children Conference* (pp. 470-474).
- Hullman, J., & Diakopoulos, N. (2011). Visualization rhetoric: Framing effects in narrative visualization. *IEEE transactions on visualization and computer graphics*, 17(12), 2231-2240.
- Hullman, J., Drucker, S., Riche, N. H., Lee, B., Fisher, D., & Adar, E. (2013). A deeper understanding of sequence in narrative visualization. *IEEE Transactions on visualization and computer graphics*, 19(12), 2406-2415.
- Kanharaj, S., Leong, R. T. K., Lin, X., Masry, A., Thakkar, M., Hoque, E., & Joty, S. (2022). Chart-to-text: A large-scale benchmark for chart summarization. arXiv preprint arXiv:2203.06486.
- Kong, H. K., Liu, Z., & Karahalios, K. (2018, April). Frames and slants in titles of visualizations on controversial topics. In *Proceedings of the 2018 CHI conference on human factors in computing systems* (pp. 1-12).
- Knaflic, C. N. (2015). *Storytelling with data: A data visualization guide for business professionals*. John Wiley & Sons.
- Kosara, R., & Mackinlay, J. (2013). Storytelling: The next step for visualization. *Computer*, 46(5), 44-50.
- Lee, B., Riche, N. H., Isenberg, P., & Carpendale, S. (2015). More than telling a story: Transforming data into visually shared stories. *IEEE computer graphics and applications*, 35(5), 84-90.
- Lundgard, A., & Satyanarayan, A. (2021). Accessible visualization via natural language descriptions: A four-level model of semantic content. *IEEE transactions on visualization and computer graphics*, 28(1), 1073-1083.
- Ma, K. L., Liao, I., Frazier, J., Hauser, H., & Kostis, H. N. (2012). Scientific storytelling using visualization. *IEEE computer graphics and applications*, 32(1), 12-19. <https://doi.org/10.1109/MCG.2012.24>
- Milligan, J. N., Hutchinson, B., Tossell, M., & Andreoli, R. (2022). *Learning Tableau 2022: Create effective data visualizations, build interactive visual analytics, and improve your data storytelling capabilities*. Packt Publishing Ltd.
- Mittal, V., Moore, J., Carenini, G., & Roth, S. F. (1998). Describing complex charts in natural language: A caption generation system. *Computational Linguistics*, 24(3), 431-477.
- Nichols, E., Gao, L., & Gomez, R. (2020, October). Collaborative storytelling with large-scale neural language models. In *Proceedings of the 13th ACM SIGGRAPH Conference on Motion, Interaction and Games* (pp. 1-10).
- Qian, X., Koh, E., Du, F., Kim, S., Chan, J., Rossi, R. A., ... & Lee, T. Y. (2021, April). Generating accurate caption units for figure captioning. In *Proceedings of the Web Conference 2021* (pp. 2792-2804).
- Qin, X., Luo, Y., Tang, N., & Li, G. (2020). Making data visualization more efficient and effective: a survey. *The VLDB Journal*, 29, 93-117.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., ... & Sutskever, I. (2021, July). Zero-shot text-to-image generation. In *International Conference on Machine Learning* (pp. 8821-8831). PMLR.
- Ren, P., Wang, Y., & Zhao, F. (2023). Re-understanding of data storytelling tools from a narrative perspective. *Visual Intelligence*, 1(1), 11.
- Roselli, D., Matthews, J., & Talagala, N. (2019, May). Managing bias in AI. In *Companion Proceedings of The 2019 World Wide Web Conference* (pp. 539-544).
- Satyanarayan, A., Moritz, D., Wongsuphasawat, K., & Heer, J. (2016). Vega-lite: A grammar of interactive graphics. *IEEE transactions on visualization and computer graphics*, 23(1), 341-350.
- Segel, E., & Heer, J. (2010). Narrative visualization: Telling stories with data. *IEEE transactions on visualization and computer graphics*, 16(6), 1139-1148.

- Stahl, B. C., & Eke, D. (2024). The ethics of ChatGPT— Exploring the ethical issues of an emerging technology. *International Journal of Information Management*, 74, 102700.
- Stokes, C., Setlur, V., Cogley, B., Satyanarayan, A., & Hearst, M. A. (2022). Striking a balance: reader takeaways and preferences when integrating text and charts. *IEEE Transactions on Visualization and Computer Graphics*, 29(1), 1233-1243.
- Srinivasan, A., Drucker, S. M., Endert, A., & Stasko, J. (2018). Augmenting visualizations with interactive data facts to facilitate interpretation and communication. *IEEE transactions on visualization and computer graphics*, 25(1), 672-681.
- Team, B. D. (2014). Bokeh: Python Library For Interactive Visualization. Bokeh Development Team.
- VanderPlas, J., Granger, B., Heer, J., Moritz, D., Wongsuphasawat, K., Satyanarayan, A., ... & Sievert, S. (2018). Altair: interactive statistical visualizations for Python. *Journal of open source software*, 3(32), 1057.
- UNESCO, C. (2021). Recommendation on the ethics of artificial intelligence.
- Ware, C. (2013) *Information Visualization: Perception for Design*; Elsevier: Amsterdam, The Netherlands; Morgan Kaufman: Boston, MA, USA.
- Wilkinson, L. (1999) *The Grammar of Graphics*; Springer: New York, NY, USA.
- Yuan, A., Coenen, A., Reif, E., & Ippolito, D. (2022, March). Wordcraft: story writing with large language models. In *27th International Conference on Intelligent User Interfaces* (pp. 841-852).
- Yonghui W. (2018). Smart compose: Using neural networks to help write emails. *Google AI Blog*.