Understanding the perceived importance of disclosing ethical concerns underlying data visualizations

Ana Carolina Ferreira Junger[1], **Simone Diniz Junqueira Barbosa** [1]

[1]Departamento de informática - PUC - Rio
Rio de Janeiro, RJ, Brazil

{ajunger, simone}@inf.puc-rio.br

Abstract. Introduction*: Data visualizations are increasingly being used as a means of conveying data. However, ethical issues of data visualizations are still an understudied topic.* ***Objective****: Our objective is to understand how ethical aspects relate to data visualization and how visualization readers perceived the importance of disclosing ethical concerns underlying data visualizations.* ***Methodology****: We reviewed the literature to understand how ethical aspects relate to data visualization and then evaluated, with visualization readers, what aspects they deem more important in two different scenarios.* ***Results****: The results show that the order of priority varies between scenarios and between participants' profiles, but understanding the target audience's needs, the visualization objective, and how the data was collected features among the most important aspects in every case. The items perceived as most important were not aligned with those found more frequently in the literature, which highlights the importance of evaluating what has been found in the literature with real participants.*
Keywords *data visualization, ethics, visualization metadata*

1. Introduction

Data visualizations have the potential to reveal our social world [Hill et al. 2016] and are the primary means of conveying information about a large data group in a simple, clear, and concise way [Wanzer et al. 2021, Pandey et al. 2014].

There are many people are involved in technology development [Coeckelbergh 2020]. In the case of data visualization, this involves collecting, processing, organizing, or displaying data. The potential impacts of data visualizations are significant; therefore, one should not neglect ethical and moral standards or fail to take responsibility [Correll 2019].

However, how ethical issues are applied in data visualizations, how they are evaluated and discussed, and how they are communicated to the readers are still subjects that remain relatively under-discussed [Nunes Vilaza et al. 2022, Carvalho et al. 2022]. Nunes Vilaza *et al.* conducted a scoping review on articles published at conferences held by the ACM Special Interest Group on Computer-Human Interaction (SIGCHI) from 2010 to 2020 and which contained, in their title, abstract, or authors' keywords list, the word "ethic*" [Nunes Vilaza et al. 2022]. Their results show a growth in interest in this subject since 2018, reinforcing the need to discuss it. However, they only got 5 (3.8% of the total 129 results) papers regarding data repositories, and only 1 [Correll 2019] regarding data visualization. Carvalho *et al.* (2022), conducted a review to understand how researchers address ethical aspects in the *IHC* conference series. Their results indicate an increase in concern from 2015 onwards; however, none of the 25 articles they selected are explicitly related to data visualization [Carvalho et al. 2022]. This underline that ethics in data visualization is still an understudied topic.

Recently, Wang *et al.* designed ethical cards to promote the debate of ethical considerations of artificial intelligence (AI) for data visualization [Wang et al. 2024]. Their work represents a significant step forward in the debate on the ethical aspects of AI techniques for data visualization (AI4VIS) technologies. They focus on promoting debate and reflection within the internal development team about the application of AI for data visualization, but they do not provide a framework to support visualization *readers* directly.

There is often a gap between communicating and understanding data [Kelley et al. 2009]. Correll (2019) states that data visualizations are designed artifacts that may represent flawed, incomplete, and potentially idiosyncratic observations [Correll 2019]. Also, graphs are commonly associated with scientific fact [Correll e Heer 2017] since humans tend to believe in numbers, take them as the purest truth and an immutable fact [Hill et al. 2016], and have difficulty recognizing uncertainty [D'Ignazio e Klein 2020]. As the main objective of visual analytics is to improve knowledge generation between humans and machines, readers must be able to trust the result [Sacha et al. 2016]. Therefore, it must find ways to support readers in understanding the data more comprehensively, including its flaws and uncertainties.

This work is a part of a bigger project aiming to define and evaluate ways of documenting design decisions that could influence the interpretation and understanding of data visualization, highlight uncertainties, elucidate ethical dilemmas that designers and developers may have contemplated, and which cannot be conveyed solely through a data visualization. In doing so, we expect to support a more comprehensive and ethical interpretation and understanding of visualizations for the reader. We describe here the initial steps: the literature study to select the content of the documentation and an evaluation of its importance by visualization readers. In this way, the main objective of this work is to evaluate what complementary information about the visualization and what decisions made by designers and developers the reader finds more important to have, accompanying the data visualization in different contexts.

Pereira *et al.* collected, in the *IHC* conference, the II Grand Research Challenges for Human-Computer Interaction (II GranDIHC-BR) [Pereira et al. 2024]. This project was a collaborative effort involving several HCI researchers from different regions of Brazil and some other countries. One of the identified challenges was to "Develop methods and techniques that promote Human-Data Interaction with Data Literacy and Usable Privacy" (GC5) [Coleti et al. 2024], which is closely aligned with our research project and the study reported in this paper.

Recently, documentation standards have been created for datasets [Gebru et al. 2021, Holland et al. 2018], machine learning (ML) models [Mitchell et al. 2019], and privacy policies [Kelley et al. 2009, Kelley et al. 2013], all with positive evaluations by designers, developers, and users. Some researchers have studied the use of metadata as documentation on data visualization

[Burns et al. 2021, Burns et al. 2024], but in a limited way (see Section 2.2). These works highlight the importance of creating a documentation model for data visualizations and evaluating how it may influence the reader's interpretation and understanding.

Researchers and data visualization creators must reflect on the power that data visualizations can have, as they increasingly shape people's understanding of our societies and our environment [Dörk et al. 2013]. It is therefore paramount to involve the people affected by the visualization [D'Ignazio e Klein 2016] and give them tools to understand the data because visualizations usually create a separation between the people impacted by the data and the people consuming the data [Correll 2019]. When creating a visualization, the designer makes several choices to represent the data. In many cases, these data have already undergone several transformations, from their collection to processing. These choices are usually not communicated to the reader, which can hinder their understanding and interpretation of the visually represented data.

We set out to evaluate what information is perceived by readers as most (or least) important in two different scenarios. To do this, we conducted a studyt inspired by Burns *et al.*[Burns et al. 2021, Burns et al. 2024], which is fully described in Section 3.

The remainder of this paper is organized as follows. Section 2 presents recent studies about data visualization ethics, the concept of annotation and metadata of data visualizations, which grounds this work, and the documentation tools that we took as references. Section 3 presents the literature review study that we conducted and the empirical evaluation study. Section 4 presents and discusses the results of the evaluation study, and Section 5 presents the concluding remarks.

2. Background and Related Work

2.1. Ethics on data visualization

Correll (2019) argues that all visualization research has a moral character, and, from there, he raises questions and dilemmas about ethics in data visualization [Correll 2019]. According to him, there is a misconception that data visualization is something neutral; however, "data are not naturally occurring phenomenon," and using data to represent people has political power [Correll 2019]. In this way, data visualization is not a neutral instrument; it is an active communication between designers and readers [Verbeek 2008, Cavanagh 2008] and a visual tool for assimilation and knowledge production [D'Ignazio e Klein 2016].

Data visualizations can persuade the reader, primarily if they do not have a strong opinion on the visualization topic [Pandey et al. 2014]. As designers significantly influence how data is interpreted [Correll e Heer 2017], and since people have some political views [Cairo 2019], it is important to have a critical eye on data visualizations because humans tend to believe in numbers and take them as the purest truth and an immutable fact [Hill et al. 2016]. However, when presented in a visualization, they are just designed artifacts, which may represent flawed, incomplete, and potentially idiosyncratic observations [Correll 2019].

Exploring the political side of visualizations Dork *et al.* (2013) proposed four principles to explore issues of power in visualizations:

- **Disclosure**, present the designer's decisions and intentions to increase trustworthiness;
- **Plurality**, present different versions to improve interpretation for each reader;
- **Contingency**, give the reader tools to conclude rather than predetermined;
- **Empowerment**, let the reader question the visualizations and use them to present their story.

We explore these principles in our work, focusing on "disclosure" by presenting the designer's decisions and intentions not only to increase trustworthiness but also to improve the readers' interpretation and understanding, and "contingency" by giving the reader more information to conclude.

D'Ignazio and Klein (2016) proposed six feminist principles for data visualization and suggested several questions for designer reflection for each principle [D'Ignazio e Klein 2016]:

- **Rethink Binaries**, not only in data visualization, but also in data collection and classification, we should consider the diversity of possibilities, not only binaries;
- **Embrace Pluralism**, explore implicit and explicit designer decisions (align with [Dörk et al. 2013]);
- **Examine Power and Aspire to Empowerment**, involve the people affected by the visualizations in the processes of creating a visualization;
- **Consider Context**, different contexts can influence different perceptions, it is always important to consider how a particular visualization output may be received;
- **Legitimize Embodiment and Affect**, acknowledging the importance of embodiment and affect while evaluating a visualization;
- **Make Labor Visible**, explicitly state the people who handle the data for each part of the process.

D'Ignazio and Klein (2016) highlight essential values that should be considered in data visualization [D'Ignazio e Klein 2016], like the importance of multidimensional gender representation, since genders are not just binary [Trans Student Educational Resources 2015]. They state that focusing on the designer's subjectivity can help them expose decisions and add to the complementarity between the designer and the reader. They also reinforce the importance of visualizing all parts of the process of creating a visualization from the very beginning of data collection. We aim to reinforce this by providing the reader with more information about data visualization.

Wang *et al.* (2024) introduced AI-VIS EthiCard, a card-based approach for designers to discuss and reflect on ethics in the relationship between AI and data visualization. The cards are organized according to thematic categories [Wang et al. 2024]:

- **Goals**, to promote the discussion of possible results and facilitate the process of conceptualizing;
- **AI-VIS Tasks**, present tasks related to AI4VIS;
- **Technologies**, presents possibles AI4VIS technology;
- **Ethical Principles**, present eleven principals based on previous researches;
- **People-In-Focus**, possible characteristics of end users;

- **Challenges**, related to ethical dilemmas.

They suggest several dynamics for the team involved in the visualization development process. We hope to contribute with this reflection, but primarily supporting the people affected by the visualization.

2.2. Annotations and metadata in data visualization

Annotations are items designers or developers add to a visualization; they can be in the graphic, textual, or social form [Hullman e Diakopoulos 2011]. They can contribute to the reader's interpretation and understanding by adding a layer of information about things that are not explicit in the core visualization and can highlight data provenance or make uncertainty explicit, for example [Hullman e Diakopoulos 2011]. In addition, they can also contribute to data recall [Borkin et al. 2016], and to the storytelling process [Segel e Heer 2010, Kong et al. 2019].

Metadata, by contrast, is defined by Burns *et al.* (2021) as "information not directly represented in a visualization that provides background on the source of the data, the transformations applied to the data, the visualization elements, its purpose, the people involved in its creation, and its intended audience" [Burns et al. 2021]. They discuss the pros and cons of metadata related to visualization pipelines and the people involved (creators and intended audience).

In the visualization pipeline, Burns *et al.* (2021) included metadata such as the data source, the process of cleaning and processing data, and the visual encoding, explaining some problems the reader could face while decoding the visualization or explaining how to read it. The pros they identified are aligned with enabling reproducibility and increasing credibility, transparency, and trust. The cons are related to privacy and trust issues, highlighting some of society's prejudices and potentially creating bias. In the metadata about people, they focus on two groups of people involved in the process: the creators, which can increase transparency and credibility, and the intended audience, which evidence the context and can help the readers to understand that not every visualization will be enlightening to everyone. The cons are related to decreased trust in the visualization if there is no trust in the creator and excluding people from understanding the visualization if they do not match the intended audience.

The decision to expose metadata to readers will always have pros and cons. Our primary goal is to understand how beneficial it will be to expose this information to the reader without the potential for causing them, people involved in the creation of the data visualization, or people whose data is exposed, significant harm.

From their work, Burns *et al.* (2024) ran two experiments to evaluate the impacts of metadata on visualizations [Burns et al. 2024]. The first evaluated what metadata people find more important or interesting for some goal. They selected a set of visualizations, and for each visualization, they asked which kind of metadata the participant preferred to see to reach a specified goal, like confidence in the chart or topic, key takeaways, assessing trustworthiness, understanding of method, design, or perspective, and satisfying interest. Their result shows that the participants were more interested in metadata related to the explanation of the visualization for various goals and information about the data source for the goal of "assessing trustworthiness."

Our work is inspired by Burns *et al.* [Burns et al. 2021, Burns et al. 2024]. However, first, we search the literature to understand how ethical aspects related to data visualization, including more information than just metadata. We did not classify pros and cons ourselves. Instead, we decided to let study participants judge what is essential and what is not important to know in different contexts. This way, we overcome a limitation that we identified in their work. This work is fully described in section 3.

2.3. Documentation tools

This section presents three documentation models and their different uses: Datasheet for datasets (subsection 2.3.1), Model Cards (subsection 2.3.2), and privacy Nutrition Labels (subsection 2.3.3). We have based ourselves on models from other areas, such as machine learning and datasets, because they have proven to be very interesting in providing a more complete view of the information in question, and we believe that a similar model for data visualization can bring many benefits to readers.

2.3.1. Datasheets for datasets

Gebru *et al.* (2021) proposed Datasheets for datasets, a dataset documentation tool [Gebru et al. 2021]. They argued that the datasets are directly related to the performance of models, and there was still no standard way of documenting datasets. They claimed that datasheets for datasets "have the potential to increase transparency and accountability within the machine learning community, mitigate unwanted societal biases in machine learning models, facilitate greater reproducibility of machine learning results, and help researchers and practitioners to select more appropriate datasets for their chosen tasks." They focus on two stakeholders: dataset creators, where the main objective is to promote reflection while creating, distributing, and maintaining a dataset; and dataset consumers, where the main objective is to ensure that they have all the necessary information to make decisions based on the dataset.

From experiences with product teams, they realized that users were more likely to answer a question about legal and ethical considerations if integrated into other relevant sections of the dataset. They, therefore, grouped the questions into the following sections:

- **Motivation**: Questions related to the motivation behind creating the dataset, people involved in the development, and if it is related to any funding interest.
- **Composition**: Questions directly related to the composition of the dataset to inform any information needed by the data consumers.
- **Collection Process**: This set of questions should elucidate the data collection process, including mechanisms or procedures used, people involved in the process, the timeframe, ethical considerations, and others.
- **Preprocessing/cleaning/labeling**: Question involving all the processes already carried out on the "raw" data.
- **Uses**: Questions related to the intended uses for the dataset and where it should not be used.
- **Distribution**: Questions related to the distribution of the dataset, like if it will be distributed internally or to third parties, how it will be distributed, copyright, and others.

- **Maintenance**: Issues related to the maintenance plan and consumer communication.

Since their launch, Microsoft, Google, and IBM have piloted datasheets for datasets in their internal teams. They highlighted that datasheets for datasets are not a perfect solution to mitigate all bias and potential risks; however, they proved to be an essential tool to facilitate the communication between data creators and consumers, and prioritize transparency and accountability.

2.3.2. Model Cards

Model cards, proposed by Mitchell *et al.* (2019), is a way of documenting ML models [Mitchell et al. 2019]. They verified that there was no standardization on model documentation, which led to errors and biases being discovered late, when the system was already in use. They also highlighted the importance of understanding the ethical considerations of models beyond just traditional metrics and statistics. The proposed solution involved the creation of documentation comprising sections that would not necessarily fit in all cases and would depend on the model, the context, and the stakeholders, and which could also include details beyond the proposed documentation, as well as datasheets for datasets. The proposed model has the following format:

- **Model Details**: basic information about the model (*e.g.*, date, version, people involved, type, license);
- **Intended Use**: "what the model should and should not be used for, and why it was created," should include the primary usage intentions, the principal intended users, and "out-of-scope uses."
- **Factors**: factors related to model performance on groups, instrumentation, and environments, expanded on relevant factors and evaluation factors.
- **Metrics**: measures used to evaluate performance, chosen decision thresholds, and approaches used for uncertainty and variability;
- **Evaluation Data**: details about the dataset, motivation, and important information about the pre-processing process;
- **Training Data**: Similar to the evaluation data, but with information about the training data, if it is possible;
- **Quantitative Analyses**: "the results of evaluating the model according to the chosen metrics, providing confidence interval values when possible."
- **Ethical Considerations**: describe ethical considerations taken in developing the model.

Nunes *et al.* conducted a speculative design study in which participants used Model Cards to help document ethics-related design decisions [Nunes et al. 2022]. An interesting finding of their work was a selective attitude regarding what to document: "positive (ethical) stances were much more likely to make it into the written document than those considered to be unethical" [Nunes et al. 2024]. Potentially negative or unethical aspects of the decisions were considered and conveyed orally, but only infrequently put in writing; and if so, participants would distance themselves from the negative consequences, *e.g.*, mentioning "the system" or adopting a passive voice structure instead of writing from a first-person perspective as the decision makers. Despite this limitation, they did acknowledge the value of using such a tool, as it helped them reflect on the ethical considerations of their designs more deeply than they would have done otherwise.

2.3.3. Nutrition Labels

Kelley *et al.* (2009) adapted the concept nutrition labels for technologies [Kelley et al. 2009]. They started from the principle that users are concerned about their privacy, but the privacy policies of websites are confusing due to the use of technical jargon [Jensen e Potts 2004]. They proposed a *Privacy Nutrition Label*, a visual representation of the main aspects of the privacy policy with non-technical text exemplifying what data is collected and how the information is used. To evaluate it, they conducted an empirical study comparing standard privacy policies and their Privacy Nutrition Label. The result shows that the participants could understand the Privacy Nutrition Label better than the standard policies.

In 2020, Apple introduced *app privacy details*,[1] Privacy Nutrition Labels for apps set on the Apple Store, so that users can understand what data is collected and how it is used. Li *et al.* (2022) evaluated the labeling process from the developer's perspective [Li et al. 2022]. Their results show that the developers agreed that the Privacy Labels benefit the users, but they found the form filling challenging. Some participants had trouble understanding it and made errors while filling out the form. However, a participant said they switched to a simpler library to collect less user data, demonstrating a positive aspect of the developer-facing possible issues with their application. Similarly, Google introduced Privacy Labels[2] in 2022 on Google Play Store.

The nutrition label format was also used to analyze and evaluate datasets. Holland *et al.* (2018) developed a framework that allows the developer to visualize a dataset before developing a model: The Dataset Nutrition Label (The Label) [Holland et al. 2018]. In 2020, Chmielinski *et al.*(2020) developed a web-based GUI interactive version with three panels: Overview, Use Cases & Alerts, and Dataset Info, where the user can explore information from the dataset to understand and manage bias [Chmielinski et al. 2020].

3. Investigating Ethical Aspects of Data Visualization

In this section, we report a literature review of ethics-related information items related to data visualization and present the study we conducted to understand which information visualization readers find more important to complement data visualizations.

3.1. Literature Review

In this section, we describe our literature review to understand how ethical aspects are related to data visualization and what information could be ethically complementary to visualization readers. We conducted a literature review following these steps: defining the research questions; creating and refining the search string; defining the inclusion and exclusion criteria; defining a quality assessment procedure; collecting data; and analyzing

[1] https://developer.apple.com/app-store/app-privacy-details/ Access date: 07/07/2025

[2] https://support.google.com/googleplay/android-developer/answer/10787469?hl=en Access date: 07/07/2025

data. Our review was inspired by [Kitchenham et al. 2009]; the main difference is that we used an automated script to assess the quality of the papers we found based on weighted terms found in the title, keyword, and abstract. In addition, we did not define a data extraction form *a priori*, to allow for flexibility in the data we selected to analyze.

Our literature review aimed to answer the following questions:

1. How are ethical aspects related to data visualization?
2. How does the context of the visualization/target audience/publishing medium and the designer's communicative intentions affect how the designer deals with ethical aspects?
3. What methods/tools exist to assist designers in considering ethical aspects of data visualization?

To create our search string, we combined terms related to both ethical aspects and data visualization. We searched on the ACM Digital Library,[3] IEEExplorer,[4], and Scopus[5], since they are the primary digital libraries in the field of computer science.[6] After some experimentation with search terms, we selected *ethics*, *ethical*, *responsibility*, *responsible*, *accountability*, *accountable*, *liability*, *liable*, *moral*, *fairness*, *transparent*, *transparency*, *bias*, and *trust*. And for terms related to data visualization, we selected *visualization*, *visualisation* (to include the American and British variations), and *visual analytics*'. The final string was *(ethics OR ethical OR responsibility OR responsible OR accountability OR accountable OR liability OR liable OR moral OR fairness OR transparent OR transparency OR bias OR trust) AND (visualisation OR visualization OR "visual analytics").* We searched within the articles' keywords because we wanted articles whose primary focus was some ethical aspect of data visualization.

To filter the articles, we adopted the following inclusion criteria:

- I1 Works that describe an ethical aspect in data visualization
- I2 Case study that analyzes some ethical aspects in data visualization systems
- I3 Studies that propose ways to deal with ethical aspects in data visualization systems
- I4 Peer-reviewed articles

And, the following exclusion criteria:

- E1 Studies that do not explicate an ethical aspect
- E2 Papers written in languages other than English or Portuguese
- E3 Studies posted before 2014, due to the growing appreciation of ethical aspects [Nunes Vilaza et al. 2022]

We performed the search, and found 1,048 articles. After excluding the duplicates, we were left with 879 articles. We applied exclusion criterion [E3] (leaving 624 papers), then [E1] (599 papers); they were all written in English or Portuguese, so [E2] did not apply. To prioritize the reading, we conducted a quality assessment step to rank the papers according to the adherence to our scope. In the interest of time, we created a script to apply

[3]https://dl.acm.org/ Access date: 07/07/2025
[4]https://ieeexplore.ieee.org/Xplore/home.jsp Access date: 07/07/2025
[5]https://www.scopus.com/ Access date: 07/07/2025
[6]We also translated the search string to Portuguese and searched in the SBC Open Library, but found no results.

scores to each paper title, keyword, abstract, and total (sum of the previously mentioned scores). Each time a relevant term appeared in the title, abstract, or keyword, points were added to the paper's score. The script was created by the first author and reviewed by the second and is available at https://github.com/carolfjunger/slr-notebook/tree/main. We acknowledge the limitation of this approach and, in future work, we intend to sample a few excluded papers to verify the effectiveness of this decision. Scores ranged from 1 to 68, with a 30 median. We selected the articles with 50+ points for full-text reading; this included 39 articles. After fully reading them to extract the content relevant to our documentation, we ended up using the information from 10 studies. We discarded 29 papers that were out of scope: 6 focused on a toolkit or framework, without addressing a concept; 4 focused on an artificial intelligence tool, not visualization; 8 did not focus on data visualization at all; and 11 did not address ethical aspects in data visualization.

For each paper, we extracted any information that could be interesting to complement the visualization. After that, we codified each item as its main topics and aggregated similar items, organizing the items in the format we hoped to have in our documentation; in this way, we obtained the first version of our documentation. We identified 24 items and 31 codes (an item could be linked to multiple codes), as illustrated in Table 1. The table is ordered according to the number of articles that mention the related item, from the most mentioned (in four articles) to the least mentioned (in one article).

Table 1. Items identified in the literature

Identified Items	Codes	Articles
What information about your black-box model can contribute to the explainability of the visualization (*e.g.*, show shap values, show the decision tree, show the distribution of features)?	explicability, transparency	[Correll 2019, Fischer et al. 2022, Yan et al. 2024, Cabrera et al. 2019]
What information can reduce the degree of uncertainty in the visualization (*e.g.*, margin of error, information on model paremetrisation, data processing uncertainty measures, possible measurement errors, input, low resolution and incomplete data, use of unofficial data, some distrust of the data source)	uncertainty	[Sacha et al. 2016, Panagiotidou et al. 2022, Correll 2019, Fischer et al. 2022]
Is there an alternative visualization design that could contribute to a more complete understanding of the visualization?	alternative visualization	[Correll 2019, Panagiotidou et al. 2022, Wang et al. 2024, Hepworth 2020]
How was data categorization represented?	data diversity, data processing, fairness	[Hepworth 2020, Cabric et al. 2024, Cabrera et al. 2019]
What are the expected impacts of visualization? And which should not happen? What are the negative consequences of visualizing data in this way and what precautions have been taken to avoid them?	people impacted, visualization goal	[Wang et al. 2024, Hepworth 2020, Correll 2019]
Are marginalized groups represented? Why? (*e.g.*, non-binary people, black people, LGBTs)	data diversity	[Cabric et al. 2024, Wang et al. 2024, Fischer et al. 2022]
What processes have been carried out on the data (*e.g.*, cleaning, aggregation, coding, dimensionality reduction)?	data processing	[Correll 2019, Wang et al. 2024, Hepworth 2020]

Continued on next page

Table 1 – continued from previous page

Identified Items	Codes	Articles
How was the data collected? (*e.g.*, logs from a tool, interview/forms, observation)	data collection	[Correll 2019, Hepworth 2020]
What is the goal of the visualization?	people impacted, target audience, visualization goal	[Hepworth 2020, Correll 2019]
Present the data source.	data source	[Wall et al. 2019, Hepworth 2020]
Which model metrics can help complete the visualization information? (*e.g.*, Accuracy, Recall, Specificity, Precision, Negative Predictive Value, False Negative Rate, False Positive Rate, False Discovery Rate, False Omission Rate, and F1 score.)	model metrics	[Yan et al. 2024, Cabrera et al. 2019]
What target audience is expected to consume the data? And which is not expected?	target audience	[Hepworth 2020, Correll 2019]
Who/which institution collected the data? (*e.g.*, ibge, google analytics, the team that made the visualization)	data collection	[Correll 2019, Hepworth 2020]
What is the appropriate reproduction medium for the context of your visualization?	reproduction medium	[Hepworth 2020, Wang et al. 2024]
What interactions can the user make?	interaction	[Fischer et al. 2022]
Who is responsible for any problems that may arise from the communication of data?	responsibility	[Fischer et al. 2022]
How were the colors chosen to represent the genres (*e.g.*, standard library options)?	color, data diversity	[Cabric et al. 2024]
Does the visualization provide guidance to help the reader understand the information?	guidance	[Fischer et al. 2022]
What is your audience's main expected need?	target audience	[Hepworth 2020]
Why is the data sorted in this way? Is there any degree of importance?	data sorting	[Hepworth 2020]
Does the data collected comply with the LGPD or GDPR?	LGPD/GDPR	[Fischer et al. 2022]
Are there any special procedures for data from marginalized groups and/or children?	data diversity	[Wang et al. 2024]
Who/which institution processed the data? (*e.g.*, ibge, third parties, the team that made the visualization)	data processing	[Correll 2019]
Present the variability of your data (*e.g.*, range, interquartile range, variance, standard deviation)	variability	[Correll 2019]

To understand the importance of each item, we ran an evaluation study with target readers, described in the following section.

3.2. Evaluation study

The primary objective of our evaluation was to understand which items were deemed more important by the readers in two distinct scenarios. We ran an online questionnaire (Google Forms), distributed on our social media network (Instagram and LinkedIn) and message groups (Whatsapp), our expected participants was anyone over the age of 18 who reads and interprets data visualizations. The questionnaire was in Portuguese, the participants' native language. The form was available for three weeks, in April 2025. Our study was inspired by Burn *et al.*'s research [Burns et al. 2021, Burns et al. 2024]. We share the goal to understand the impact of metadata in data visualization. However, from the literature review, we decided to include more information items than pure metadata. We asked participants, in the role of visualization readers, to evaluate the importance of different information items.

3.2.1. Questionnaire structure

The questionnaire was divided into the following sections:

1. A small introduction contextualizing the study
2. The Informed Consent Form, describing the general and specific objectives of the evaluation, the risks and benefits, the procedure, the ethical considerations, and ensures confidentiality and anonymity. All participants could abandon the questionnaire whenever they wanted without penalties. This evaluation study is part of a research project approved by the Research Ethics Committee at PUC-Rio, protocol 97/2020.
3. Basic info about the participant, like age range, highest education degree, and primary graduation area.
4. Questions about the participant's expertise in data visualization. We asked, on average, how often they read and interpret graphs and how often they design data visualizations, from never to every week. We also asked their level of expertise in the chart types we included in our study (simple bar chart, line chart, stacked bar chart, grouped bar chart, bubble chart, and boxplot), from : none to expert.
5. Question to pseudo-randomize the order of the scenarios. We asked whether their birthday was an even or odd number. If the participant selected odd, they would see the hospital scenario before the movie scenario (*HM*group). Otherwise, they would see the movie scenario before the hospital scenario (*MH*group).
6. For each scenario, the questionnaire had the following sections:
   (a) Explanation of the scenario and two visualizations as examples.
   (b) Items classification: We asked them to score the items we identified in the literature review on a 7-point scale, from “Not important at all" to “Essential". We added the option "I do not know how to classify it" and an open text field if the participant wanted to add an item that they thought was missing
   (c) Top 5 items: We asked them to select the 5 most important items for the scenario.
   (d) Bottom 5 items: We ask them to select the 5 least important items for the scenario.
7. And, in the final section:
   (a) An open question that the participant could add anything they want.
   (b) For those who would be willing to deepen the subject in a one-on-one interview, we asked for their contact information, explaining that this info would not be linked to their responses in the form.

For each scenario, we decide to ask the participant to classify the items by importance and to select the top and bottom 5, since the items were mapped from the literature, we hypothesized that they would classify them in a high level of importance, and by the next step (to selected the top and bottom 5), we can get the most and least important.

3.2.2. Questionnaire content

We aimed to evaluate the items' importance in two different scenarios, one more sensitive and one more mundane. For this, we defined a hospital and a movie scenario. We presented the scenario in a different order for each participant group so we could evaluate whether it would affect the participants' responses.

In the hospital (*H*) scenario, the participant should play the role of a policymaker deciding where to build a new hospital based on two visualizations with information about the imaginary city: a boxplot with the age of the residents of each cuty region, and a bubble chart with the quantity of hospitals by specialty and region.

In the movie (*M*) scenario, the participant should play the role of a worker at a movie company that should select which genres the public is most interested in to organize next year's release schedule based on two visualization: a grouped and stacked bar chart with the movie classification by genre, star, and month, together with the number of active users per month; and a line chart depicting the percentage error from a machine learning model, with predictions from the social media network of the users' engagement with a publication per month, per gender. The visualizations used in each scenario can be accessed at https://carolfjunger.github.io/form1/.

We pilot-tested the questionnaire, one of the authors accompanied the pilot participant while the participant answered the questionnaire. We asked the participant to narrate their steps and ask questions when necessary. They took 32 minutes to finish it. Based on the participant's feedback, we improved its structure. In each section, we present the scenario description and visualizations. While piloting, we noticed that the participant read the whole description again between the item classification and the top and bottom sections, and was disappointed when they realized it was the same text. To mitigate this burden, in the top and bottom sections, we explained that we were just repeating the scenario description and visualizations in case the participant wanted to read it again.

4. Results and Discussion

We received 24 responses. However, one participant said that they never read or interpret data visualizations and declared that they do not have any familiarity with the chart examples we provided. Also, they very often answered with "I do not know how to classify it", so we decided to remove this participant from the analysis as not meeting the intended profile. 11 participants started the questionnaire with the movie scenario (henceforth *MH* group) and 12 participants with the hospital scenario (henceforth *HM* group). Most (19) participants were 26 to 39 years old; two were 40 to 59 years old, and two were over 60 years old. Most (20) participants have or are currently studying for a postgraduate degree; the distribution of education levels is shown in Table 2.

We divided the analysis of their educational background into two groups: STEM+ – Science, Technology, Engineering, and Mathematics (STEM), and Economics, which involve more knowledge of mathematics and data visualization – and non-STEM, comprising all other areas. We had 10 STEM+ and 13 non-STEM participants. Regarding the participants' knowledge of data visualization, the majority read or interpret graphs every week (9 participants) or every month (9); the other 3 read or interpret every semester, and 2 every year. Concerning designing visualizations, 10 participants never design visualizations, 1 designs every year, 5 every semester, 5 every month, and 2 every week. As illustrated in Figure 1, all participants have some level of knowledge of bar and line charts. Most participants have from moderate to high knowledge of bar charts (19 participants), line charts (19), and stacked bar charts (17). About the grouped bar chart,

Table 2. Participants' Education Level

Education Level	Count
University degree (incomplete)	1
University degree (complete)	2
Post-graduation/specialization (incomplete)	1
Post-graduation/specialization (complete)	4
Master's degree (incomplete)	3
Master's degree (complete)	4
Doctor's degree (incomplete)	4
Doctor's degree (complete)	4

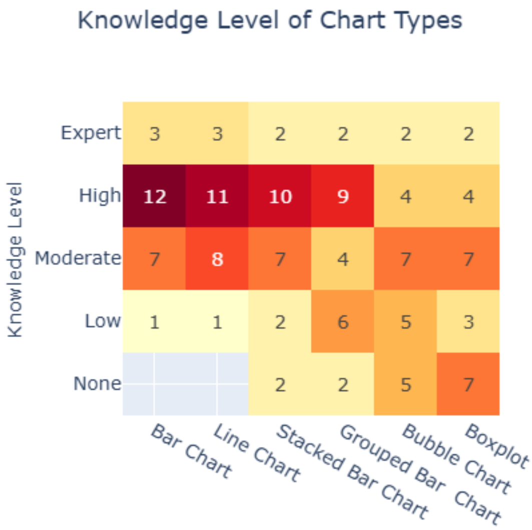

Figure 1. Knowledge level by chart Type

the biggest groups are divided between high (9) and low (6) knowledge. And, for bubble chart and box plot, we have the highest number of participants who have no knowledge (5 and 7, respectively). We note that the questionnaire tool (Google Forms) we adopted did not allow us to add an example image for each chart, so participants needed to know the chart names.

4.1. Item Classification

In the items classification phase, the participant should classify each item from "Not important at all" to "Essential" in a 7-point scale; they also had the option "I do not know how to classify it." We do not consider empty values in our analysis; as these cases were infrequent, they did not significantly interfere with the results. We carried out various analyses: overall results, by each scenario, by the order in which each participant saw the scenario, and by the participants' background area. In none of the analyses was the average of any item extreme, either positive or negative. This was to be expected, as we used a 7-point scale [Joshi et al. 2015]. To classify each item, we used the average rating as described in [Pimentel 2019].

Overall, the five items with the highest (top) ratings were: *guidance*, *target audience needs*, *objective*, *how data is collected*, and *uncertainty*, all of them classified, on average, as *Important*. The five items with the lowest (bottom) ratings were: *reproduction medium*, *explicability*, *responsibility*, *data sorting*, and *color selection*. The first four were classified, on average, as *Neutral*, and *color selection* was classified, on average, as *Not important*. The remaining items were classified, on average, as *Slightly important*, as exemplified in Figure 2.

In the hospital scenario, the top-rated items were *target audience needs*, *uncertainty*, *objective*, *how data is collected*, and *guidance*, all classified, on average, as *Important*; *data source* was the top sixth, with the same classification. And the five

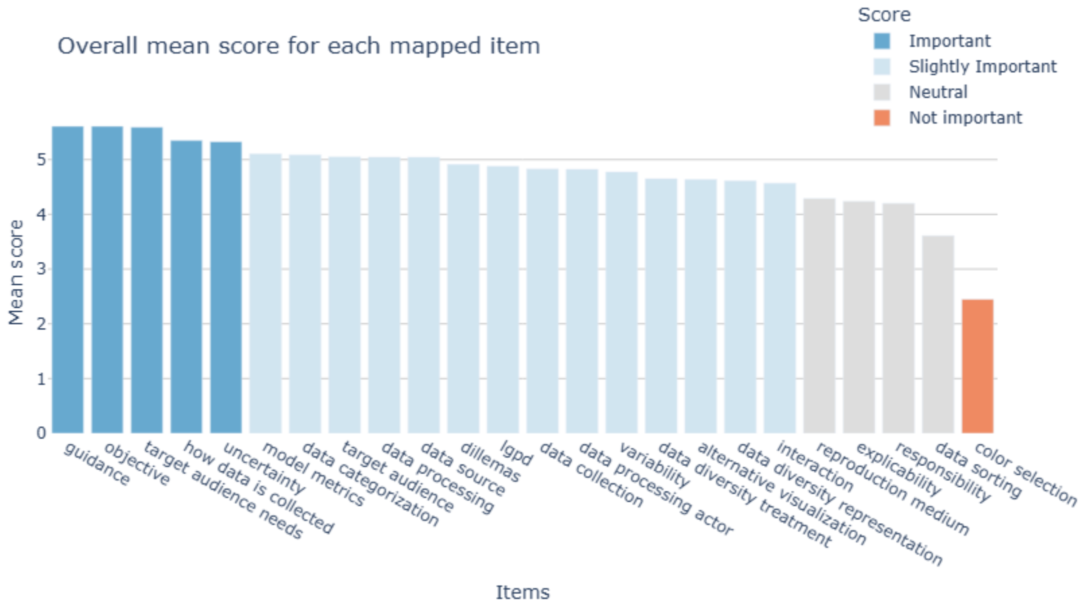

Figure 2. Overall classification

bottom-rated were *data diversity treatment* (classified, on average, as *Slightly Important*), *data diversity representation*, *responsibility*, *reproduction medium*, and *data sorting* (classified, on average, as *Neutral*). We did not expect the low ratings of the *data diversity treatment* and *data diversity representation* items, as it is important to draw attention to marginalized groups' problems and needs [Cabric et al. 2024] and because *data diversity representation* were mentioned in three articles in our review (the most mentioned item were mentioned in four articles). However, none of the visualizations that we exemplified had a gender or race-specific cut; we hypothesize that this could be one of the reasons for the low score in this scenario. The score rankings are represented in Figure 3.

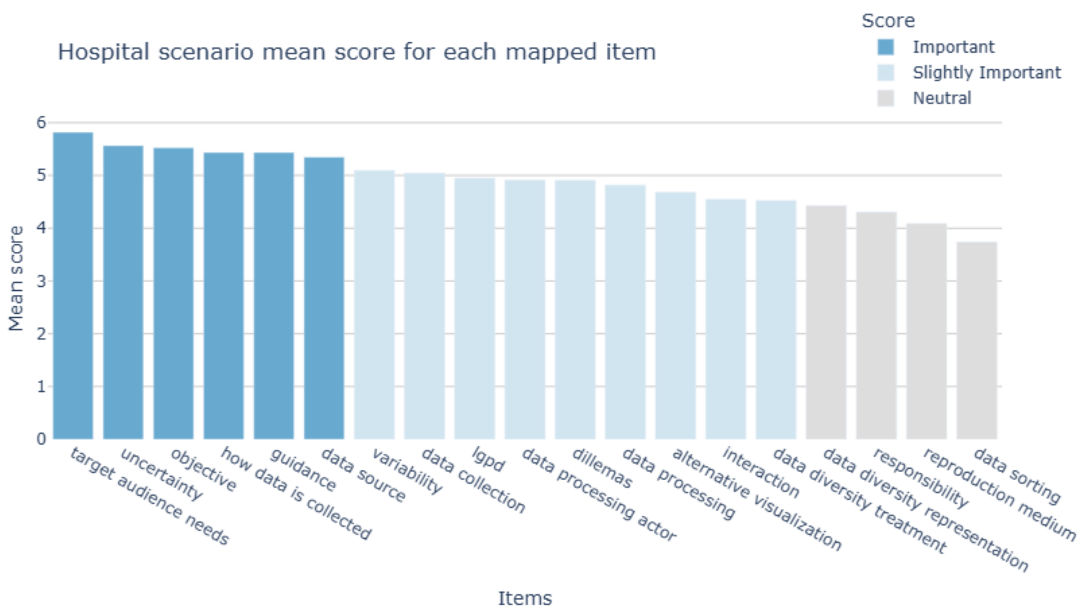

Figure 3. Hospital scenario classification

In the movie scenario, the top five rating items were *guidance*, *objective*, *target audience needs* (all three classified as *Important*), *data processing*, and *how data is collected* (both classified as *Slightly Important*). The least five items were variability (classified as *Slightly Important*), *explicability*, *responsibility* (classified as *Neutral*), *data sorting* (classified as *Slightly Not Important*), and *color selection* (classified as *Not Important*). In this scenario, the participants classified the items on average slightly positively; we hypothesize that a reason could be that it was a more mundane scenario, so the items would end up being just a complement to the visualization, not something extremely necessary, but not compromising either (Figure 4).

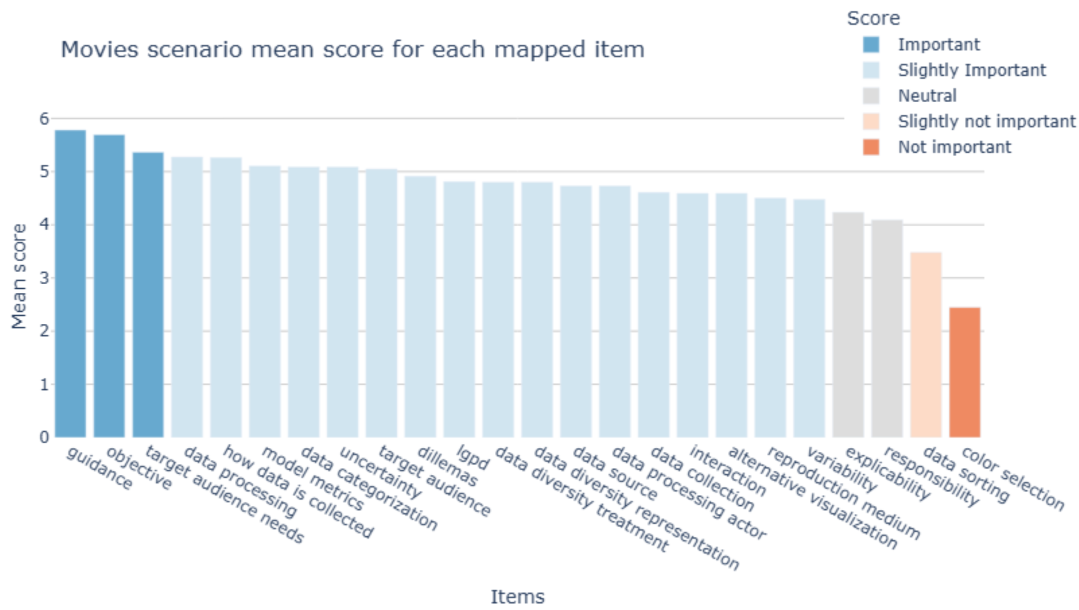

Figure 4. Movie scenario classification

The four most important ones are the same in both scenarios: *target audience needs*, *objective*, *how data is collected*, and *guidance*. And, *responsibility* and *data sorting* are the least important items in common. It is worth highlighting that *explicability*, *color selection*, and *model metrics* did not appear as options in the hospital scenario, because were items related to a genre cutoff or related to model results, that did not appear in the visualizations of the hospital scenario. In this sense, in the movie scenario, where we present visualization with a genre cutoff, the items related to data diversity became more important; however, they did not score in the top 5.

The scenario order influenced the preferences of the participants a little. *HM* participants only classified as *Important*, on average: *guidance*, *target audience needs*, and *objective*. *MH* participants classified, on average, nine items as *Important*. As we can see, the order of the scenarios impacted the participants' decisions regarding the most important items: the participants who saw the more sensitive scenario before the mundane one gave lower scores overall. The scenario order did not have an effect for the items judged as least important: *color selection*, *data sorting*, and *explicability*.

Analyzing the participants' primary background, the STEM+ participants selected more items with a high importance score than participants from other backgrounds. The participants with STEM+ graduation classified, on average, eleven items as *Important*, and for non-STEM, only four items were classified, on average, as *Important*. For participants with primary graduation areas of STEM+, the most important item is *model metrics*, which may be related to a greater knowledge of data science, since three participants in other areas selected the option *I do not know how to classify it* for *model metrics*; however, as we did not ask about their knowledge of data science, we cannot state that. We can also notice that for STEM+ participants, *data diversity representation* is *Neutral* and one of the least important items, for non-STEM+ participants, *data diversity representation* is classified as *Slightly Important*. One possible reason for this is that non-STEM people, especially in the humanities, are more likely to engage in debates about diversity and have a more diverse cultural background.

4.2. Top and Bottom Five Items

In this section, we analyze and discuss the results from the step of selecting the top and bottom five important items. In this step of the questionnaire, the participant should select the five most and least important items to accompany the visualization for each scenario. In this section of the questionnaire, the items were presented to the participants randomly, so the order of the items did not bias the results.

Overall, the items that were selected most often as most important were *target audience needs*, *objective*, *how data is collected*, *guidance*, and *data processing* (Figure 5). The most often selected as least important were *data sorting*, *responsibility*, *reproduction medium*, *LGPD*, and *color selection*. All items were selected as least important at least once, and *data sorting* and *color selection* were the only items that were not selected as most important not once.

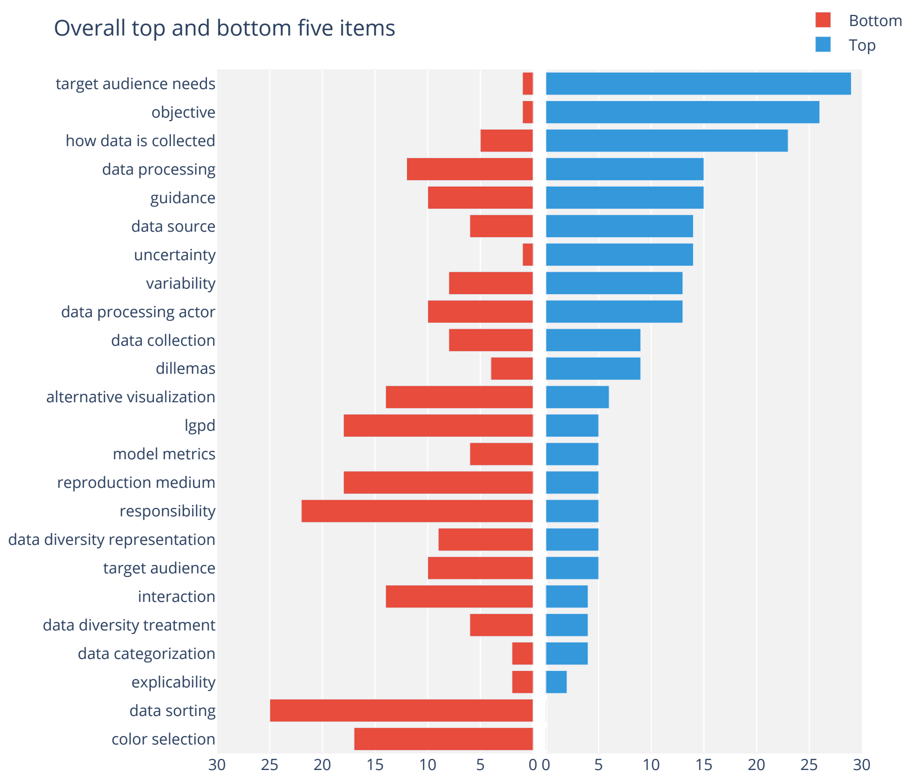

Figure 5. Overall result for top and bottom five items

For the hospital scenario, the most frequent top items were *target audience needs*, *objective*, and *how data is collected*. *Target audience needs* were selected as the top important for 20 of 23 participants (almost 87%) . We expected *guidance* to appear more prominently, because it was the context with the boxplot and bubble chart, the two graphs that participants said they knew the least about, and the *guidance* prompt text was *"Include guidance to help the reader understand the information."* Based on this result, we hypothesize that although the participants said they did not have much knowledge of this type of graph, they were able to understand it intuitively, so *guidance* was not very important. The most selected as least important were *data sorting*, *reproduction medium*, *LGPD*, *target audience*, and *interaction*. The results are exemplified in Figure 6.

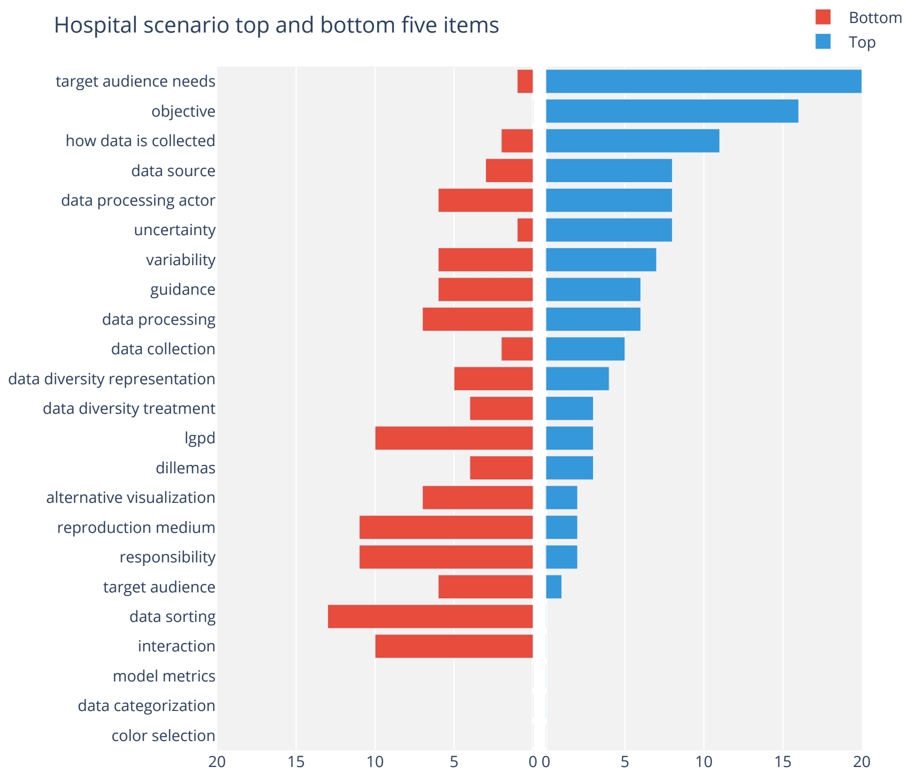

Figure 6. Hospital scenario top and bottom five items

For the movie scenario, the most frequent top items were *how data is collected*, *objective*, *target audience needs*, *guidance*, and *data processing* (Figure 7). Of these, only the *target audience needs* was not selected as one of the least important, along with *dilemmas* and *uncertainty*. And the most selected as least important were *color selection*, *data sorting*, *responsibility*, *LGPD*, *reproduction medium*, and *alternative visualization*. In this scenario, we underline that although there were graphs in which the participants had more knowledge, when we added more information together, and increased the dimension of the graph, the grouped stack bar plot that has four dimensions (movie genres, star, month, and user) and the sum of all users at the top, the participant find more important to have some *guidance*, in contrast with a graph that he knows less about, but with fewer dimensions, the case of the Hospital scenario.

Analyzing the influence of the order of the scenarios, for the participants in the *HM* group had disagreement between the participants about *data processing*, it was marked 9

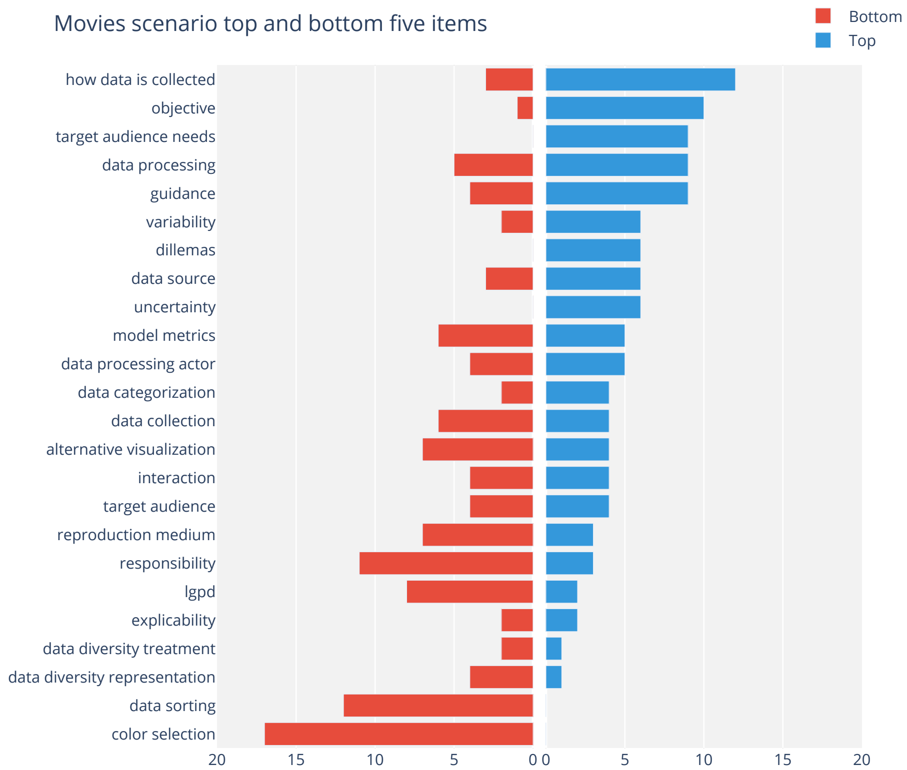

Figure 7. Movie scenario top and bottom five items

times as most important and 8 times as least important, in which case one participant marked *data processing* as top and least for the same scenario, we hypothesize that could be a consequence of fatigue, which should be investigated in future studies. In contrast, the participants in the *MH* group agreed more, since for the four most selected as top items, only one was selected as the bottom item, and only once. The complete result is shown in Figures 8 and 9.

Analyzing by background, participants from both areas (STEM+ and non-STEM+) classified the most important items as *target audience needs*, *objective*, and *how data is collected*. The STEM+ participants also classified *data processing* as one of the top five important items. However, *data collection* was classified as one of the least important items. For these participants, it is more important to know how the data was collected and processed than who collected the data; we hypothesize that this can be related directly to the fact that *responsibility* was also one of the least important, since most STEM+ participants design visualizations (9 of 10 STEM+ participants). As such, they may have put themselves in the designer's role and felt motivated to ascribe responsibility to the "system" instead of themselves [Barbosa et al. 2021]. This is an important but subtle issue that should be further explored in the future. The complete result is exemplified in Figures 10 and 11.

4.3. Study Limitations

The number of respondents was far lower than Burns *et al.*'s (2024) study [Burns et al. 2024], we got 23 valid answers, and they had 64. We share a limitation, most of our participants were highly educated. Also, most (19) of them were younger

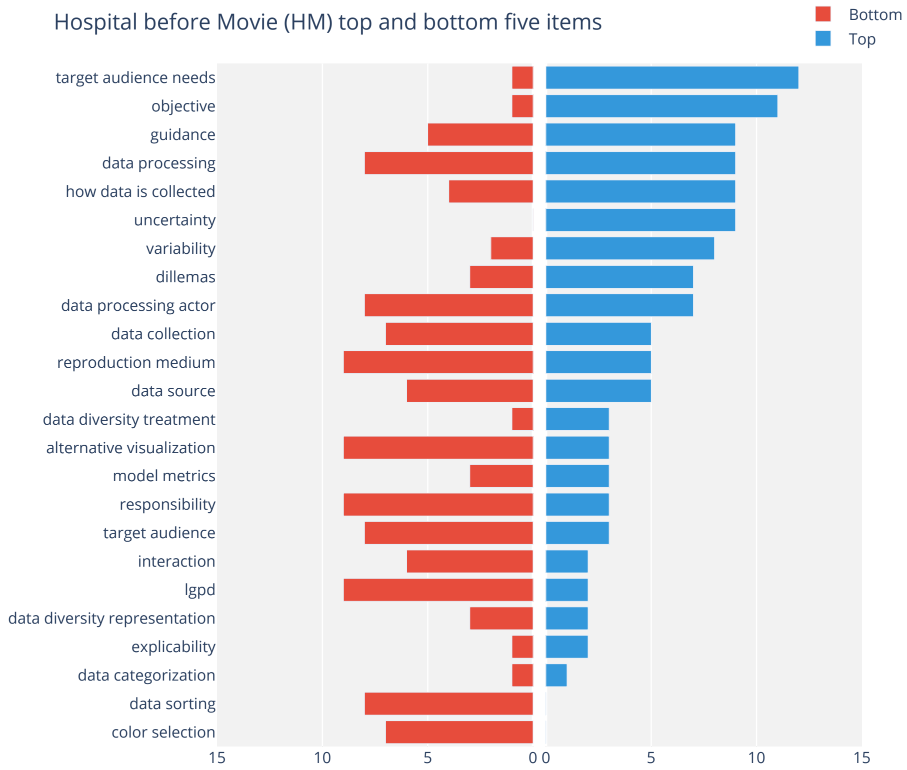

Figure 8. *HM* group top and bottom five items

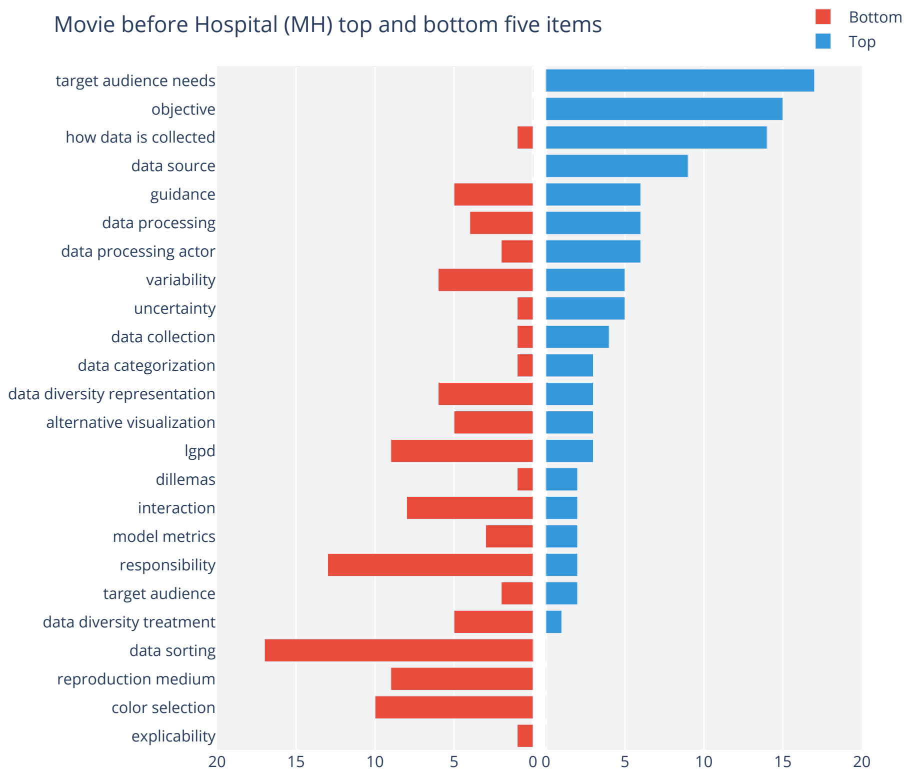

Figure 9. *MH* group top and bottom five items

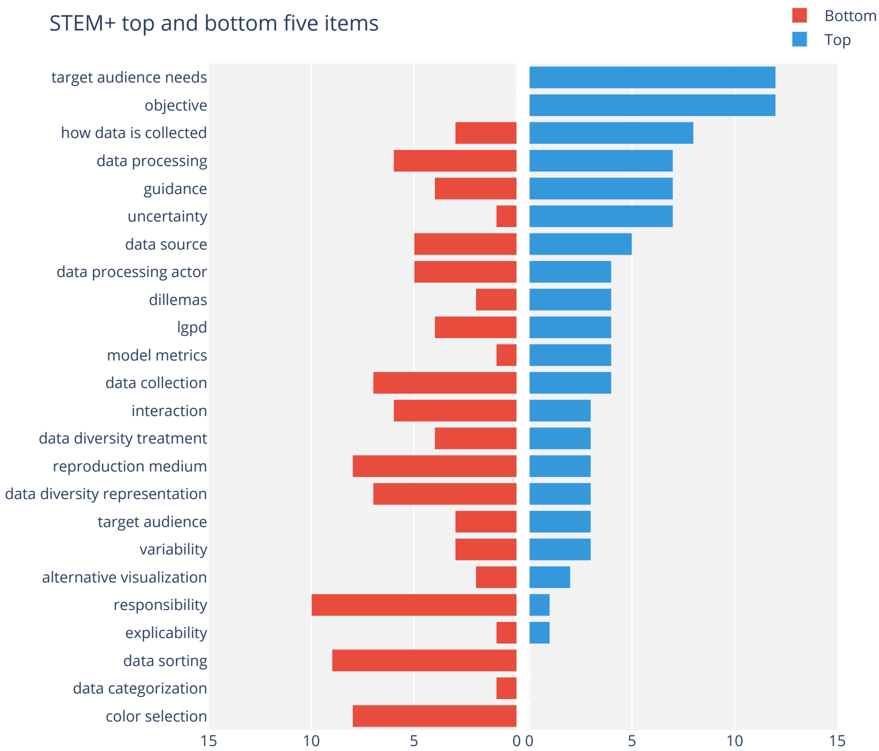

Figure 10. STEM+ top and bottom five items

adults (26 to 39 years old), and most (20) are highly educated (have or are currently studying for a postgraduate degree), which does not represent data visualization readers in general, however, despite their high level of education, some (10 participants) had low to none knowledge of, at least, one chart types that we used. In future studies, we hope to improve the distribution of our questionnaire in order to reach a wider and diverse audience, including non-expert users.

Another limitation was that some items only appear in one scenario: *explicability*, *color selection*, and *model metrics* only appear in the movie scenario, which could cause some bias in the result. Also, since we gave concrete examples to exemplify the item in each scenario, the chosen example may have influenced the participant's decision.

Some participants declared they had no knowledge of some of the charts we had selected, which may have influenced their decision about what is important or not. We also used some terms that the participants may not know about, for example, in the movie scenario, one visualization is the result of a machine learning model, which may be unfamiliar to some participants. However, as we did not ask about their knowledge of machine learning models, we cannot draw any conclusions about this.

5. Conclusion and future works

With this experiment, we learned a few lessons. Which items are important for data visualization readers is a very personal decision that varies from individual to individual and from context to context, but there are some items that remain important or unimportant in most cases. Understanding the *target audience's needs*, the visualization *objective*, and *how the data was collected* was deemed paramount in the identified

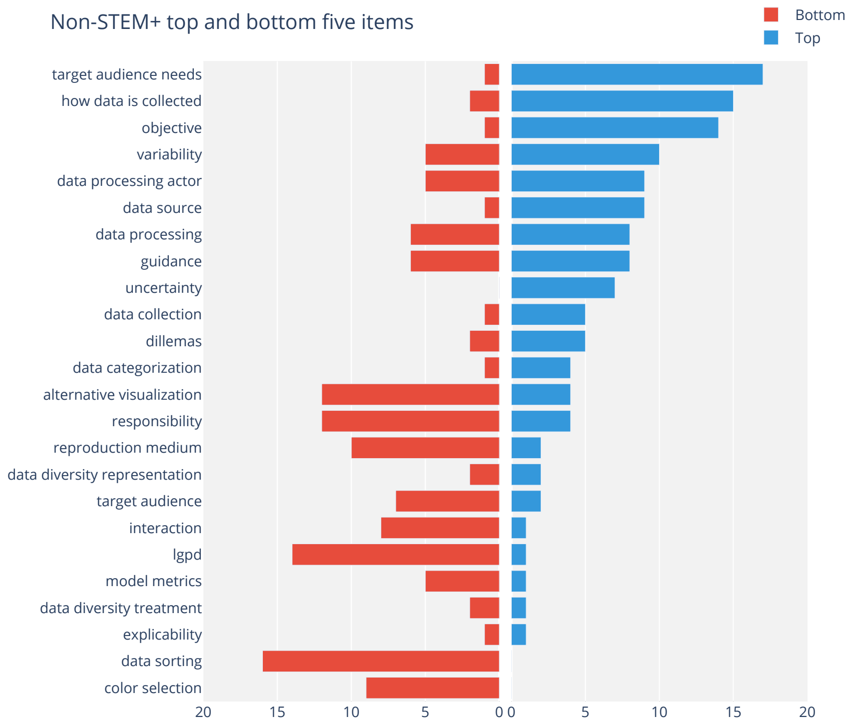

Figure 11. Non-STEM+ top and bottom five items

scenarios and for most readers. These were the only items evaluated as important in the overall items classification and selected the top 5 items in both scenarios[7]. Moreover, how colors were chosen to represent genders (*color selection*), how the data was sorted (*data sorting*) in the visualization, knowing the *reproduction medium* expected for the visualization, and existing a responsible for possible problems with the visualization (*responsibility*) are information that the participants did not find important to have. These items were classified on average as *Neutral* or *Not important* (only *color selection*) and at the bottom 5 in at least one scenario. [8]

One participant suggested items that we had not identified in the literature: allowing different colors for accessibility purposes. We also add that, beyond physical disabilities, visualization designers need to aim at providing access to information for people with Intellectual and Developmental Disabilities (IDD) [Wu et al. 2021], as well.

The participants' responses were not closely aligned to the findings in the literature. We thought they would classify most items found in the literature as *Essential* or *Important*. However, they deemed several items as *Not important at all* or *Not important*. This results alert us that the literature in visualization ethics cannot be taken as direct guidelines and highlight the importance of evaluating existing, published results with real participants, in studies situated in the specific context at hand. As opposed to our hypothesis, in the overall results, most of the items (14) were classified on average as *Slightly Important*, even though it is a positive classification, we expected that more items

[7]Since there were items only present in one scenario, for the top and bottom classification, we analyze by the top 5 of each scenario, as in the item classification we used the mean of the score

[8]All items, excepted for *Color selection* that was not evaluated in the hospital scenario, appear in the bottom 5 in both scenario.

would be classified at a high rate. The items classified as important make interceptions with the top 5 items from the next step, as mentioned in the previous paragraph. Also, we expected that *guidance* would be perceived as more important where the participants had less knowledge of the chart; however, this was not the case.

Another interesting factor that we mapped was whether the visualization followed the LGPD or GDPR, which was considered *Important* or *Slightly Important* in almost all analyses (only non-STEM+ participants mapped it, on average, as *Neutral*). However, it was chosen as the least important item in most analyses. Surprisingly, *data diversity representation* and *data diversity treatment* did not appear as one of the most important items in any of the contexts and scenarios analyzed, primarily for STEM+ participants, despite appearing in three articles in our literature review. For instance, [Coleti et al. 2024] consider them crucial aspects in data visualizations and claim that visualization designers have a responsibility to consider various genders, ethnicities, ages, and physical and mental abilities in order to avoid bias.

As [Huey et al. 2023] discovered, people tend to prefer visualization with minimal information. This may explain some of our findings, since we propose to make available extra information to complement the visualization. Also, visualization is commonly associated with saving time and understanding a big group of data fast [Chen et al. 2014], which is aligned with [Burns et al. 2024], which finds that people are more interested in metadata related to the explanation of the visualization. In other words, people tend to prefer something more straightforward and direct, instead of taking the time to understand and process it. Which brings us to an ethical dilemma: how should data visualization designers simplify their graphics to make it easier for the reader without losing relevant information or creating an obstacle to adequately interpreting the data?

As for future work, some (5) participants left their personal contact details, so we hope to interview them to find out more about their opinions in an open conversation. We hope that, with an open conversation, we can get the reasons underlying the participants' evaluations and better test the hypotheses of our analysis. Also, we hope to conduct real-life studies in different contexts to better understand the underlying issues. Based on the results of these studies, we plan to develop a documentation tool to accompany data visualization. We hope to contribute to the challenge (GC5) by creating a tool that facilitates users' understanding and visualization of complex data [Coleti et al. 2024].

6. Ethical Considerations

In the questionnaire, all participants agreed with the Informed Consent Form, that describe the general and specific objectives of the evaluation, the risks and benefits, the procedure, the ethical considerations, and ensures confidentiality and anonymity. All participants could abandon the questionnaire whenever they wanted without penalties. This evaluation study is part of a research project approved by the Research Ethics Committee at PUC-Rio, protocol 97/2020. All data were analyzed anonymously, and all results were aggregated into categories in which it was not possible to identify the participants.

Acknowledgements

The authors thank CAPES and CNPq (process #313049/2021-1) for the financial support for this work. We thank the anonymous participants of the study for their invaluable time.

References

Barbosa, S. D. J., Barbosa, G. D. J., Souza, C. S. D., e Leitão, C. F. (2021). A Semiotics-based epistemic tool to reason about ethical issues in digital technology design and development. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 363–374, Virtual Event Canada. ACM.

Borkin, M. A., Bylinskii, Z., Kim, N. W., Bainbridge, C. M., Yeh, C. S., Borkin, D., Pfister, H., e Oliva, A. (2016). Beyond Memorability: Visualization Recognition and Recall. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):519–528.

Burns, A., Lee, C., On, T., Xiong, C., Peck, E., e Mahyar, N. (2024). From Invisible to Visible: Impacts of Metadata in Communicative Data Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 30(7):3427–3443.

Burns, A., On, T., Lee, C., Shapiro, R., Xiong, C., e Mahyar, N. (2021). Making the Invisible Visible: Risks and Benefits of Disclosing Metadata in Visualization. In *2021 IEEE Workshop on Visualization for Social Good (VIS4Good)*, pages 11–15, New Orleans, LA, USA. IEEE.

Cabrera, A. A., Epperson, W., Hohman, F., Kahng, M., Morgenstern, J., e Chau, D. H. (2019). FAIRVIS: Visual Analytics for Discovering Intersectional Bias in Machine Learning. In *2019 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 46–56, Vancouver, BC, Canada. IEEE.

Cabric, F., Bjarnadóttir, M. V., Ling, M., Rafnsdóttir, G. L., e Isenberg, P. (2024). Eleven years of gender data visualization: A step towards more inclusive gender representation. *IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS*, 30(1):10.

Cairo, A. (2019). *How Charts Lie: Getting Smarter about Visual Information*. W. W. Norton & Company, New York, first edition edition.

Carvalho, L. P., Suzano, J. A., Santoro, F. M., e Oliveira, J. (2022). A meta-scientific broad panorama of ethical aspects in the Brazilian IHC. *Journal on Interactive Systems*, 13(1):105–126.

Cavanagh, T. (2008). Diverse designing: Sorting out function and intention in artifacts. In *Philosophy and Design: From Engineering to Architecture*, pages 301–315. Springer Netherlands, Dordrecht.

Chen, M., Floridi, L., e Borgo, R. (2014). What Is Visualization Really For? In Floridi, L. e Illari, P., editors, *The Philosophy of Information Quality*, volume 358, pages 75–93. Springer International Publishing, Cham.

Chmielinski, K., Newman, S., Taylor, M., Joseph, J., Thomas, K., Yurkofsky, J., e Qiu, Y. C. (2020). The dataset nutrition label (2nd gen): Leveraging context to mitigate harms in artificial intelligence. In *NeurIPS Workshop on Dataset Curation and Security*.

Coeckelbergh, M. (2020). Artificial Intelligence, Responsibility Attribution, and a Relational Justification of Explainability. *Science and Engineering Ethics*, 26(4):2051–2068.

Coleti, T. A., Divino, S. B. S., Salgado, A. D. L., Zacarias, R. O., Saraiva, J. D. A. G., Gonçalves, D. A., Morandini, M., e Santos, R. P. D. (2024). GranDIHC-BR 2025-2035 - GC5 - Human-Data Interaction Data Literacy and Usable Privacy. In *Proceedings of the XXIII Brazilian Symposium on Human Factors in Computing Systems*, pages 1–24, Brasília DF Brazil. ACM.

Correll, M. (2019). Ethical Dimensions of Visualization Research. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13.

Correll, M. e Heer, J. (2017). Black Hat Visualization. In *Workshop on Dealing with Cognitive Biases in Visualisations (DECISIVe), IEEE VIS, 2017*.

D'Ignazio, C. e Klein, L. F. (2016). Feminist Data Visualization. In *Workshop on Visualization for the Digital Humanities at IEEE VIS Conference 2016*.

D'Ignazio, C. e Klein, L. F. (2020). *Data Feminism*. The MIT Press.

Dörk, M., Feng, P., Collins, C., e Carpendale, S. (2013). Critical InfoVis: Exploring the politics of visualization. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, pages 2189–2198, Paris France. ACM.

Fischer, M. T., Hirsbrunner, S. D., Jentner, W., Miller, M., Keim, D. A., e Helm, P. (2022). Promoting Ethical Awareness in Communication Analysis: Investigating Potentials and Limits of Visual Analytics for Intelligence Applications. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 877–889, Seoul Republic of Korea. ACM.

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., e Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12):86–92.

Hepworth, K. J. (2020). Make Me Care: Ethical Visualization for Impact in the Sciences and Data Sciences. In Marcus, A. e Rosenzweig, E., editors, *Design, User Experience, and Usability. Interaction Design*, volume 12200, pages 385–404. Springer International Publishing, Cham.

Hill, R. L., Kennedy, H., e Gerrard, Y. (2016). Visualizing Junk: Big Data Visualizations and the Need for Feminist Data Studies. *Journal of Communication Inquiry*, 40(4):331–350.

Holland, S., Hosny, A., Newman, S., Joseph, J., e Chmielinski, K. (2018). The dataset nutrition label: A framework to drive higher data quality standards.

Huey, H., Oey, L. A., Lloyd, H. S., e Fan, J. E. (2023). How do communicative goals guide which data visualizations people think are effective? *Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)*.

Hullman, J. e Diakopoulos, N. (2011). Visualization Rhetoric: Framing Effects in Narrative Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2231–2240.

Jensen, C. e Potts, C. (2004). Privacy policies as decision-making tools: An evaluation of online privacy notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Chi '04, pages 471–478, New York, NY, USA. Association for Computing Machinery.

Joshi, A., Kale, S., Chandel, S., e Pal, D. (2015). Likert Scale: Explored and Explained. *British Journal of Applied Science & Technology*, 7(4):396–403.

Kelley, P. G., Bresee, J., Cranor, L. F., e Reeder, R. W. (2009). A "nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, Mountain View California USA. ACM.

Kelley, P. G., Cranor, L. F., e Sadeh, N. (2013). Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3393–3402, Paris France. ACM.

Kitchenham, B., Pearl Brereton, O., Budgen, D., Turner, M., Bailey, J., e Linkman, S. (2009). Systematic literature reviews in software engineering – A systematic literature review. *Information and Software Technology*, 51(1):7–15.

Kong, H.-K., Liu, Z., e Karahalios, K. (2019). Trust and Recall of Information across Varying Degrees of Title-Visualization Misalignment. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, Glasgow Scotland Uk. ACM.

Li, T., Reiman, K., Agarwal, Y., Cranor, L. F., e Hong, J. I. (2022). Understanding Challenges for Developers to Create Accurate Privacy Nutrition Labels. In *CHI Conference on Human Factors in Computing Systems*, pages 1–24, New Orleans LA USA. ACM.

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., e Gebru, T. (2019). Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 220–229, Atlanta GA USA. ACM.

Nunes, J. L., Barbosa, G. D. J., De Souza, C. S., e Barbosa, S. D. J. (2024). Using Model Cards for ethical reflection on machine learning models: An interview-based study. *Journal on Interactive Systems*, 15(1):1–19.

Nunes, J. L., Barbosa, G. D. J., de Souza, C. S., Lopes, H., e Barbosa, S. D. J. (2022). Using model cards for ethical reflection: A qualitative exploration. In *Proceedings of the 21st Brazilian Symposium on Human Factors in Computing Systems*, IHC '22, pages 1–11, New York, NY, USA. Association for Computing Machinery.

Nunes Vilaza, G., Doherty, K., McCashin, D., Coyle, D., Bardram, J., e Barry, M. (2022). A Scoping Review of Ethics Across SIGCHI. In *Designing Interactive Systems Conference*, pages 137–154, Virtual Event Australia. ACM.

Panagiotidou, G., Poblome, J., Aerts, J., e Vande Moere, A. (2022). Designing a Data Visualisation for Interdisciplinary Scientists. How to Transparently Convey Data Frictions? *Computer Supported Cooperative Work (CSCW)*, 31(4):633–667.

Pandey, A. V., Manivannan, A., Nov, O., Satterthwaite, M., e Bertini, E. (2014). The Persuasive Power of Data Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2211–2220.

Pereira, R., Darin, T., e Silveira, M. S. (2024). GranDIHC-BR: Grand Research Challenges in Human-Computer Interaction in Brazil for 2025-2035. In *Proceedings* *of the XXIII Brazilian Symposium on Human Factors in Computing Systems*, pages 1–24, Brasília DF Brazil. ACM.

Pimentel, J. L. (2019). Some Biases in Likert Scaling Usage and its Correction. *International Journal of Sciences*, 45(1).

Sacha, D., Senaratne, H., Kwon, B. C., Ellis, G., e Keim, D. A. (2016). The Role of Uncertainty, Awareness, and Trust in Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):240–249.

Segel, E. e Heer, J. (2010). Narrative Visualization: Telling Stories with Data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148.

Trans Student Educational Resources (2015). The gender unicorn. Accessed: 2024-08-31.

Verbeek, P.-P. (2008). Morality in design: Design ethics and the morality of technological artifacts. In *Philosophy and Design: From Engineering to Architecture*, pages 91–103. Springer Netherlands, Dordrecht.

Wall, E., Blaha, L., Paul, C., e Endert, A. (2019). A Formative Study of Interactive Bias Metrics in Visual Analytics Using Anchoring Bias. In Lamas, D., Loizides, F., Nacke, L., Petrie, H., Winckler, M., e Zaphiris, P., editors, *Human-Computer Interaction – INTERACT 2019*, volume 11747, pages 555–575. Springer International Publishing, Cham.

Wang, Z., Hao, S., e Carpendale, S. (2024). Card-based approach to engage exploring ethics in AI for data visualization. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems*, Chi Ea '24, New York, NY, USA. Association for Computing Machinery.

Wanzer, D. L., Azzam, T., Jones, N. D., e Skousen, D. (2021). The role of titles in enhancing data visualization. *Evaluation and Program Planning*, 84:101896.

Wu, K., Petersen, E., Ahmad, T., Burlinson, D., Tanis, S., e Szafir, D. A. (2021). Understanding Data Accessibility for People with Intellectual and Developmental Disabilities. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, Yokohama Japan. ACM.

Yan, X., Zhou, Y., Mishra, A., Mishra, H., e Wang, B. (2024). Exploring Visualization for Fairness in AI Education. In *2024 IEEE 17th Pacific Visualization Conference (PacificVis)*, pages 1–10, Tokyo, Japan. IEEE.