Fundamentals of POC, Interactive Visual Data Analysis (IVDA) Group (2024)
J. Bernard (Professor)

How to be Human-Exposing Data Humanism in Research

Roberto Honegger[1]

[1]University of Zurich, Switzerland, 16-715-419

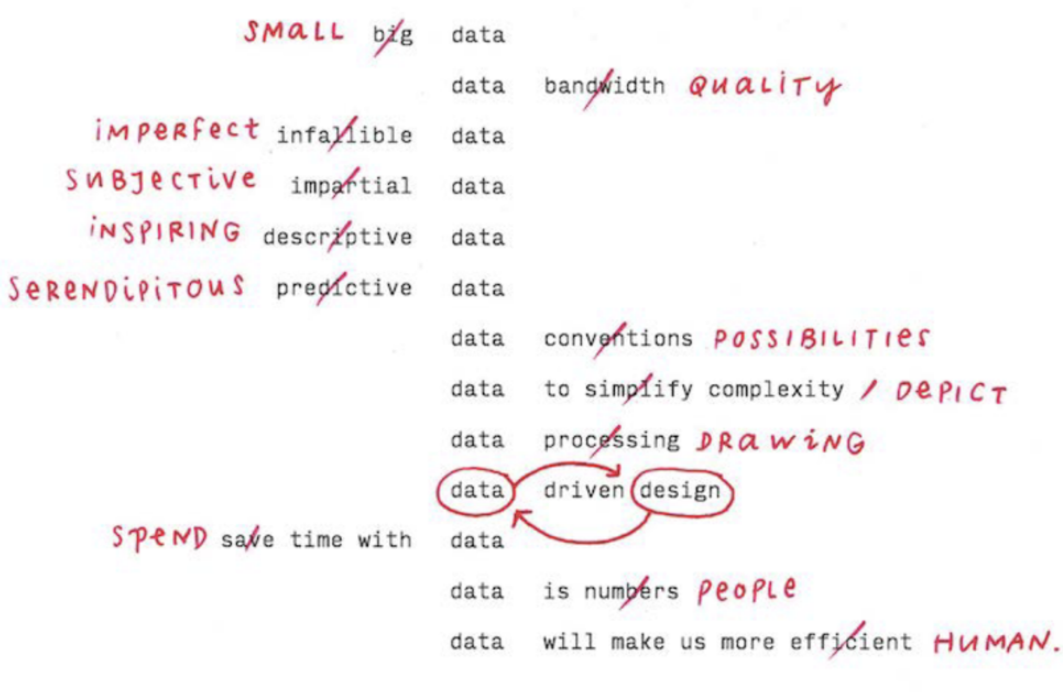

Figure 1: *We've reached peak infographics. Are you ready for what comes next? from [Lupi, 2017]*

Abstract

This paper examines the status quo of "Data Humanism" in today's Computer Science research landscape. By embedding a sample pool of papers into a layer onion-model and categorizing them along the dimensions "Storytelling," "Emotion and Empathy," "Algorithmic Self," and "Contextual Sensitivity," which are inspired by the manifesto by Giorgia Lupi, this report gives insights, whether "Data Humanism" is being reflected in recent academic papers. The chosen categories resemble the core and essence of being human and act as a method or framework to analyze whether papers adhere to the pillars of "Data Humanism." The result of this analysis is a graph-network, displaying and modeling the relationships within the sample pool of papers.

1. Introduction

Our daily lives have changed drastically with the rise of machines. Almost every little action we take involves some algorithm computing an output, whether we look for a particular restaurant to it out on Google Maps or ask our digital assistant to turn on the light. With that, we leave a vast number of digital traces called data. Nevertheless, in most cases, those data points do not capture the whole reality but rather a specific point in time about a specific focus point. In Visual Analytics, we use those data points to make sense of that data to gain some insights. For that, we rely on graphics and grammar. Good grammar allows us to gain exciting and life-changing insights and provides us with a strong foundation for a range of graphics [Wickham, 2010].

While we have made many advances in defining this grammar and conceptualizing how statistical graphics should look, we should have paid more attention to the fact that data points are digital traces made by humans. Each person has a unique story to tell, and those are not accordingly depicted in a statistical graphic. We are on the verge of questioning the impersonality of this technical approach to data. There are already initiatives trying to design new ways of connecting numbers to what they resemble: people, behaviors, and knowledge. Data represents real life. Like a picture, it is a snapshot of the world and catches a tiny moment in our lives [Lupi, 2017].

Therefore, we need to put more weight on the narrative of the data. Data Humanism seeks precisely that purpose and is a philosophy and movement started by Giorgia Lupi, an information designer and artist. Give meaning to Data to make it more human.

We see that shift in how we collaborate with Machines and Machine Learning Models. The human is not in the loop anymore, but rather, the human is in the loop. "Human in the loop" represents that human experts consult analytic algorithms occasionally to give feedback and course correction. Here, the emphasis is on identifying the work processes of analysts and integrating analytics into existing interaction processes. For instance, the sensemaking loop for analysts model, which can be see in Figure 2, was developed by Pirolli and Card. It explains the complicated interactive procedure that analysts carry out. The two main sub-loops are synthesizing hypotheses and searching for pertinent information.

According to this theory, algorithms must be completely rewritten to fit this paradigm. They must also learn from analysts' current interactions during the sensemaking process and display results in a way that makes sense within that process [Endert et al., 2014]. We see advances where algorithmic models align with human values and beliefs. This is important since algorithms predict high-stakes settings like healthcare, criminal justice, or finance. [Wang et al., 2022]

This paper examines whether scientific papers, articles, etc., have started implementing "Data Humanism" as a pillar for their work.

2. Methodology

We were given 24 papers and had to look for five additional papers, which we then had to categorize according to different parameters. Accompanying this papers, we received a template for an excel list with columns like ID, Authors, Journal, Application Domain, etc. We also had to rate the papers for the emphasis on Visual Analytics, Machine Learning, Interactive Machine Learning, and Interactive Visualizations. This already gave a specific guideline and scaffolding for categorizing the papers.

In order to structure this project and streamline the process, I

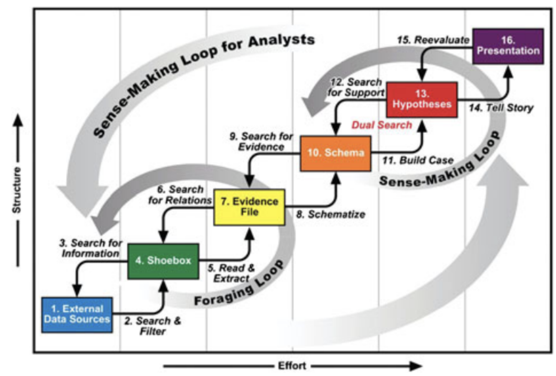

Figure 2: *The sense-making loop for analysts. From [Endert et al., 2014]*

used different tools. First, I used "Microsoft OneNote" to make notes and keep track of the requirements for this project. Then, the main workspace, "Obsidian," allows creating a graph network based on tags. So, as a starting point, I created a folder where I added all the PDFs of the report. In order to get an overview, I started reading the papers one by one and marking the most important passages. In order to highlight this, I used a community plug-in offered by Obsidian called "Annotator". "Annotator" increased my productivity by allowing me to add specific tags to each paper. So, I started categorizing each paper with the initial columns of the reference list. Furthermore, the annotator allowed me to add Hashtags to each paper, which helped me create a graph network that visualized the relationship between the papers. After each time of categorizing the reading paper, I transferred the won information into the existing reference file.

After reading all the papers and having a broad overview of the papers in the initial collection, I started thinking about additional categories for the reference list. For that, I used Giorgia Lupi's article "Data Humanism: The Revolutionary Future of Data Visualization" as inspiration. Having this text as a reference, I tried to get to the essence of being human, and I came up with "Storytelling", "Emotions&Emphathy", "Contextual Sensitivity", "Algorithmic Self", and "Cognitive Load." For each category, I created a scale to rate the papers according to Table 1.

Emotion & Empathy (EE)	Algorithmic Self (AS)	Storytelling (ST)	Contextual Sensitivity (CS)	Cognitive Load (CL)
Ignored	Not reflected	Absent	Ignored	High Load
Noted	Superficially recognized	Acknowledged	Noted	Moderate Load
Considered	Moderately modeled	Applied	Considered	Low Load
Emphasized	Fully integrated	Integrated	Emphasized	Optimized

Table 1: *The papers have been categorized by the above structure*

For those categories, I used the Onion-layer model, which is used widely in work & organizational psychology to ensure that the data point is embedded in a specific space within the individual, groups, and the overall society. As you might notice, I started with the inner layer, "Emotion&Empathy". Within the same layer, I asked the question of whether the data point is part of an "Algorithmic Self" that is being constructed. I defined the algorithmic self as the model of our own personality, which is being built within the digital space.

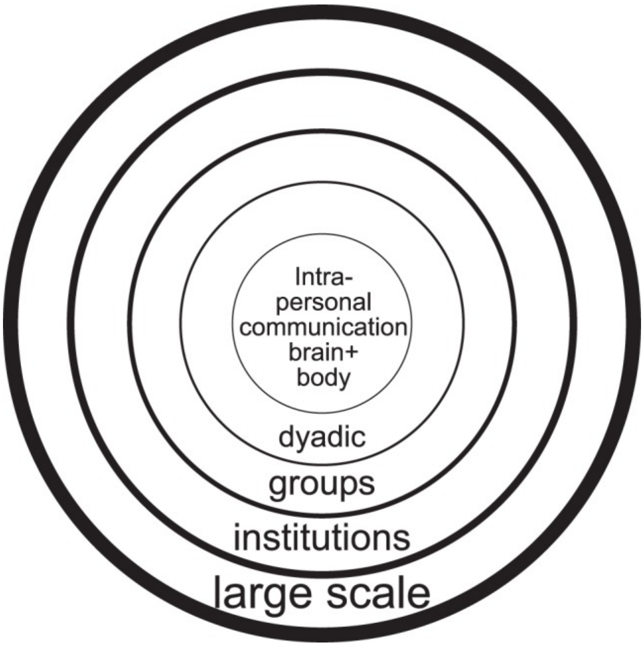

Figure 3: *Onion-Layer model of our society. From [Kappas, 2013]*

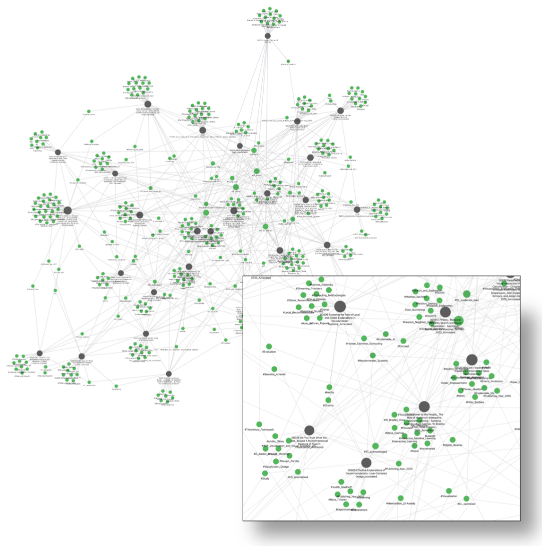

Figure 4: *The created graph network using the tags of the papers. Grey nodes depict the papers in the sample pool, while green nodes represent their tags.*

Figure 3 shows the original model, which I slightly adapted for the categories in this report.

The next category "Storytelling" is one of the most important factors of being human. No other animal uses their imagination and expressiveness to tell an abstract story. In this report, this category answers the question of whether the paper is using narrative techniques to explain the data visualization and each individual data point, meaning giving “life” to the data point and not only showing their value within a two-dimensional space, for example.

Contextual Sensitivity answers the question of whether this data point takes cultural differences, social factors, and other surrounding data into account. Therefore it surrounds the core "Emotion&Empathy", followed by "Algorithmic Self" and then "Storytelling".

The papers and the presented tools were also analyzed by their cognitive load, which refers to the amount of information our working memory can process at a specific point in time. This category was not used for the synthesis of the papers but is still mentioned by reason of completeness.

After having categorized all the papers, I started researching for new additional papers that conveyed the narrative that I wanted to create in this report. I used a Breadth-first Search algorithm, examining potential new papers adhering to the ranking table that was provided to us students. I started with IEEE Visual Analytics Science and Technology and visited their webpage and looked for recent papers. I filtered the webpage by looking at keywords like "Emotion", "Human", "Social", etc. Then, I went to Google Scholar to see whether the pdf was available to download. If yes, I downloaded it and put it into a folder, "Additional Papers". I did the same procedure for IEEE Conference on Advances in Visualization and Visual Analytics, whereby I skimmed through each year, filtering the same expressions as before. I found twenty additional papers, from which I also created a selection of the five most suitable papers that I wanted to include in this report. For this selection, I also made sure that only newer papers were included to ensure relevancy.

Succeeding the creation of the selection, I categorized them according to the structure I mentioned beforehand, and I could start the analytical and synthesis process where I could examine the papers and their tags in the following graph network consisting of nodes and edges. To create this graph, I used Obsidian’s built-in "graph-view" function. The grey nodes represent the papers, while the green nodes represent the tags that I used. For the tags, I always used the same procedure: Category + Scale = EE_ignored =Emotion and empathy ignored.

The public repository for the Obsidian Vault can be accessed under following link and can be used to recreate the network graph that I am mentioning throughout this report: https://github.com/Blockrobo/data-humanism.git

3. Synthesis of the Papers of the Reading List

Using this node graph, I was able to gain new insights about this collection of papers. Having categorized all the papers according to the tags, I could group them by filtering them by a tag or keyword. I want to share this insight, following one layer after the other as proposed by the onion model. Hereby, I used the layer "Emotion&Empathy" as the "Intrapersonal communication", "Algorithmic Self" as the "Body" layer, "Storytelling" as the dyadic layer and "Contextual Sensitivity" as the "Group" layer. As you can see,

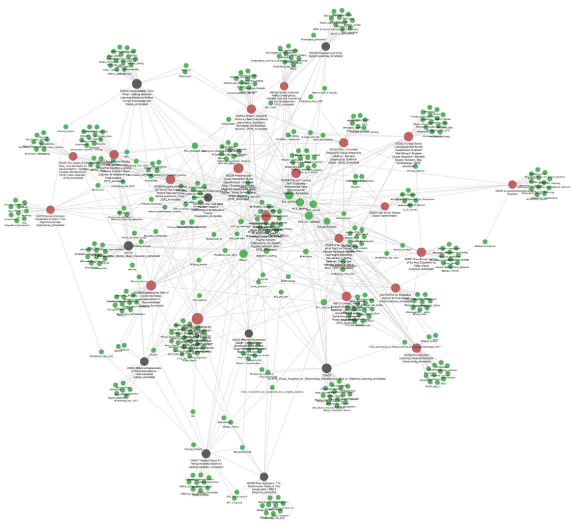

Figure 5: *Highlighting the tag "EE_ignored" results in most of the nodes turning red.*

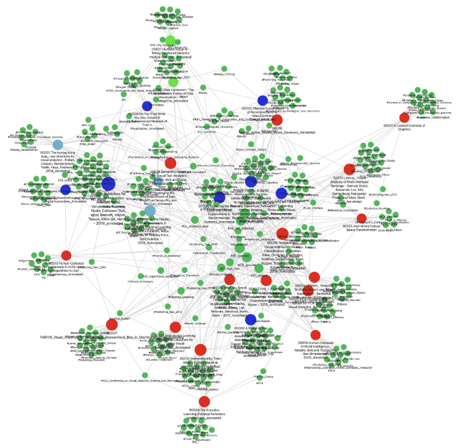

Figure 6: *Analyzing the colorcoded "Storytelling" category. Red=absent, royal-blue =acknowledged, sky-blue=applied, green=integrated*

I split the core layer from the original model into two separate layers for this paper.

Emotion&Empathy: I put the spotlight first on the category "Emotion&Empathy" by filtering the tag "EE_ignored", which describes that Emotion and Empathy were not mentioned or acknowledged at all in the specific paper.

What I observed is that a majority of the papers do not acknowledge the role of human emotion when talking about data. As we can see visually, most of the previous grey nodes became nodes. Spoken in numbers 68.965% of this sample of papers did get the tag “EE_ignored”. According to the American Psychological Association (APA), affect is any experience of feeling or emotion and is a crucial aspect of human intellect. There is a broad agreement that emotion plays a critical role in perception, cognition, and behavior and is an essential component of human intelligence. Therefore, we need to put more weight on the emotions in Visual Analytics. [Lan et al., 2023]. Xingyu Lan et al. coined the term “Affective Visualization” in their paper. They also analyzed and discussed how emotion in data visualization is underappreciated and emphasized how crucial it is to take human emotion into account when designing visualizations.

Algorithmic Self: Moving to the next layer or category, we get to the formation of an "Algorithmic Self". Two nodes light up if we look for the tag "AS_fully_integrated". These papers were the manifesto by Giorgia Lupi and the paper "Humane Visual AI: Telling the Stories Behind a Medical Condition". I want to shed light on this particular paper later in this report because it is an interesting and special one. Nonetheless, one paper regarding Recommender Systems stood out in particular, when talking about the "Algorithmic Self". In their paper "Exploring the Role of Local and Global Explanations in Recommender Systems", Marissa Radensky et al. explored what kind of influence local or global explanations have on the user. They superficially touched upon the "Algorithmic self" by modeling individual recommendations and personalization tailored to a user, thus portraying his preferences and behavior. Nevertheless, the question arises whether the system is the system really portraying us or whether it is telling us how to behave and what to like. This question is utterly important to keep in the back of our minds. [Radensky et al., 2022] Errors made by these recommender systems typically have fewer severe repercussions; they might even provide users with intriguing recommendations for films, books, or dining establishments. Malicious actors can, however, manipulate these networks to sway voting decisions, disseminate hate speech, and alter public perceptions of issues like gun control, vaccinations and climate change. Carefully considered designs that enhance user control could boost customer satisfaction and reduce harmful use. [Shneiderman, 2020]

Storytelling: As we can see, entertainment plays an important role in being human. The core of books, films, musicals, etc., is the stories that are being told and, thus, the emotions that are being transmitted from the cast to the audience. Or from one person to another. As we can see, the line between dyads and groups in the Onion Model is blurry when we talk about storytelling since it’s at the core of being human. If we examine the pool of papers in our sample, we see that only two integrate storytelling in their visualization and the narrative of the data. These are, again, the manifesto of Griogia Lupi and an additional paper, which I want to talk about later. “A Visual Active Learning System for the Assessment of Patient Well-Being in Prostate Cancer Research» by Jürgen Bernard et al. acknowledges storytelling within their tool [Bernard et al., 2015b]. This paper visualizes the electronic health care data of a patient along a timeline, such that practitioners can read the story of a person and not only their symptoms.

But if we look at the overall picture of the graph, we see that a little bit more than 55% of the papers do not include narrative ele-

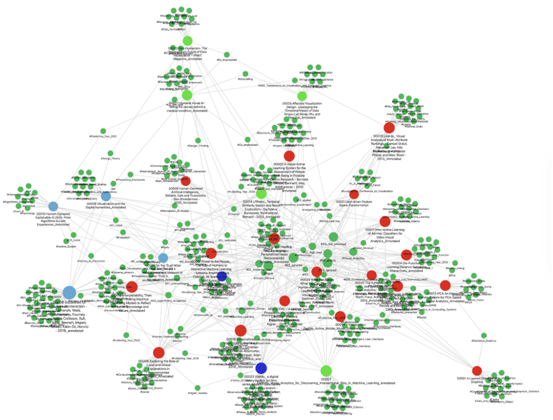

Figure 7: *Analyzing the colorcoded Contextual Sensitivity Category. Red = Ignored, royal-blue =noted, sky-blue=considered, green=emphasized*

ments to explain the origin or the circumstances of the data. They just use them as mere "fuel" for their models without taking into consideration that the data is a snapshot of human behavior directed by emotions.

Contextual Sensitivity: We are social animals, and our personality is embedded in a broader society and culture, which gives us guidelines and a moral compass on how to interact with each other. This interaction is also talked about in the paper Guidelines for Human-AI Interaction by Microsoft [Amershi et al., 2019]. The clear distinction is, though, that the paper talks not about the interaction of us humans with each other but about the interaction we maintain with AI systems. This paper postulates 18 guidelines for the Human-AI interaction, and one of those guidelines includes the “mitigation of social biases”. Under the line, it means making sure the words and actions of the AI system do not promote unfair and undesired prejudices and biases. Since our world consists of different cultures and human entities with different values and belief systems, this guideline is one of the most important ones for digital humanism. AI systems need to understand that human $\neq$ human. Each of us has a different belief system with sets of truths influenced by our surroundings or culture.

If we look at the tags for “Contextual Sensitivity”, we see that around 62% of this sample of papers ignore Contextual sensitivity and, therefore, omit the important fact of whether their models and tools can be used by different entities. There might be a social bias or any other problems that might arise by using those tools and models. One tool examines the fairness of a given visual analytical solution. FAIRVIS is a mixed-initiative visual analytics solution that enables users to audit the fairness of machine learning models by integrating a novel subgroup-finding technique. Users of FAIRVIS can generate and examine known subgroups, as well as investigate suggested and related subgroups, by using domain knowledge [Cabrera et al., 2019] . There are papers that emphasize and explore the setting of where data was sourced or where the models are applied. E.g., the papers "The Effective Explanations of Recommendations: User-Centered Design", [Tintarev and Masthoff, 2007], "LFPeers: Temporal Similarity Search and Result Exploration" [Sachdeva et al., 2023] and "VisInfo: a digital library system for time series research data based on exploratory search—a user-centered design approach" [Bernard et al., 2015a]

Trust: By respecting the core of the human essence (Emotion, Empathy, Storytelling, Contextual Sensitivity) when building visual Analytical Solutions and AI systems, we enhance the trustworthiness of the depicted data for decision-making. Trust is one of the most central topics, even for humans. Although it is a vague idea, trust is essential to human-machine interactions.

Users may be reluctant to act upon the information displayed in a visualization or to depend on it if they believe it to be unreliable. Similar to visualization, users’ trust in the technology has a major role in how well AI is integrated into businesses [Pandey et al., 2023].

Studies show that in order to build trust, the more transparently the algorithm works, the better the output is being perceived. As was proven in the lab study of TagFlip, the mechanism made it clear to users how the system worked and resulted in high user perceived transparency [Kamalzadeh et al., 2016].

Tools: As of now, AI systems and programs are tools for humanity to increase our productivity and gain more knowledge. As technology advances and AI models get smarter and we both live alongside each other, it is important that this new "creature" understands the essence of being human. Many tools that we use or propose are based on pure numbers and figures. Papers like “Dis-Function: Learning Distance Functions Interactively» [Brown et al., 2012], "LineUp: Visual Analysis of Multi-Attribute Rankings" [Gratzl et al., 2013], "iPCA: An Interactive System for PCA-based Visual Analytics" [Jeong et al., 2009] or "Semantic Interaction for Visual Text Analytics" [Endert et al., 2012] show that a lot of papers still rely on statistics and numbers without putting weight on the narrative of the gathered and used data points. Some of them try to get down to the human essence by introducing an interactive learning system where humans give feedback and readjust the model, but at the core, it is still not human. There are some other tools and papers that provide a great range of interactivity in their visualization and conception of data, like "Personalized Visual-Interactive Music Classification" [Ritter et al., 2018], "Power to the People: The Role of Humans in Interactive Machine Learning" [Amershi et al., 2014], "Inter-Active Learning of Ad-Hoc Classifiers for Video Visual Analytics" [Höferlin et al., 2012].

Nevertheless, there are advances where we can observe that emotion and empathy play a crucial role within the papers. Papers like "Human-Centered Explainable AI (XAI): From Algorithms to User Experiences" [Liao and Varshney, 2021], "User-driven Feature Space Transformation" [Mamani et al., 2013], "What You See Is What You Can Change: Human-Centered Machine Learning By Interactive Visualization" [Sacha et al., 2017], "VIAL: a unified process for visual interactive labeling" [Bernard et al., 2018]. Those papers show that we’re at the forefront of a paradigm shift where we put emphasis on the human aspect of data.

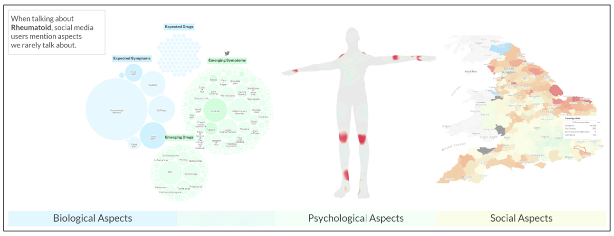

Figure 8: *Narrating the story behind a medical condition through biological, psychological, and socio-environmental lens (i.e., through the bio-psycho-social model). From [So et al., 2020]*

4. IVDA for Data Humanism

Data Humanism is a revolution in the way we see and interact with data and intelligent systems. Instead of viewing data as only numbers and figures, we give those data points life by incorporating the story and a holistic view. Many might argue that a lot of that information is unnecessary, redundant, and leads to chart junk, but this leads to a far more important question sitting at the depth of our society [Lan et al., 2023]. Nine publications of this paper argue that the present trend in visualization design to minimize emotion is not a set norm but rather a product of societal values, culture, and historical context. Emotional emphasis or repression is a sociocultural construct dominated by the humanities.

The work of the humanities needs to be brought to new life through technology, and on the other hand, advances in computer science need to be mixed with questions from the humanities to create a new epistemology that offers new avenues for knowledge. We need to start thinking of the humanities as collaborators in a new hybrid epistemology rather than as a supplier of data [Bradley et al., 2018].

Data Humanism is an important paradigm change, especially for the future of our society. We want to keep our human values while having AI systems and programs running our daily lives. Hereby, Storytelling is one of the most important points, as the paper “Do you trust what you see suggests” [Pandey et al., 2023]. Its findings imply that news media visualizations are more reliable than those created by government or scientific organizations. We need to make sure that the user understands what he or she sees to make up his own mind without being stuck in an echo chamber. This is shown in the paper “Why am I reading this” [Arnórsson et al., 2023]. Their user-friendly and comprehensible machine learning methodology assists users in comprehending, evaluating, and optimizing the personalized recommendation system.

The papers mentioned in this section represent the additional papers from this research. These papers put emphasis on the emotions, storytelling, and interactive Visualizations part. Last but not least, the paper “Human Visual AI” is the perfect example and the way ML-applications should be structured for the age of “Data Humanisms” [So et al., 2020]. Human Visual AI incorporates the biopsychosocial model from psychology and embeds every data point in storytelling so that healthcare practitioners can examine the pain a patient has to endure.

5. Discussion, Challenges, and Future Work

Given the data and information gathered from this sample of papers, data humanism is not an approach used in most papers. We see advances where Emotion, Empathy, Storytelling, and Social settings, the core of being human, become increasingly important and are also emphasized by specific research papers. Nonetheless, there has yet to be a coherent framework to analyze whether a study, tool, etc., adheres to the pillars of Data Humanism. This report is generally subjective since it only conveys my point of view and opinion. The categories I proposed have yet to be empirically tested or examined to see whether they are true. For this reason, the operationalization might be inaccurate. Several participants could replicate and fill out the study to adequately assess the documents based on the proposed categories to categorize the papers correctly. This can be a potential future research project as a meta-tool to analyze studies, papers, tools, and AI systems and systematically break down whether they adhere to the pillars of Data Humanism. Since data is a mere reflection of our thoughts, attitudes, behaviors, and opinions, it is flawed like we are [Lupi, 2017]. It might be the greatest challenge of all to create an adequate representation of what it means to be human, such that Data fed into Models, systems, and programs resembles the whole of humanity and not only a snapshot. But after all, if we convey the essence of being human: Storytelling, Emotions, and Social Interactions, we make the system understand us better and, therefore, enhance the interaction with this future potential "creature".

6. Conclusion

This paper examined whether “Data Humanism” has become a status quo in interactive data visualization research. What we conclude from our sample of papers is that there are advances in the right direction to make data more human and more embedded in our daily lives, but there is still a lot of potentials since most of the papers mostly portray statistical algorithms and computations. By embedding data into a conceptual layer model, we gain even more insights about a person and can, therefore, optimize the interaction and all the visualizations to make the person-machine relationship feel more human.

7. Use of AI Tools

I wrote this paper without the use of generative AI Tools like "ChatGPT" or "Bard". This means this paper and the synthesis of insights was solely the work of a human being. Nonetheless, "Grammarly" was used to make sure that the text was written correctly and in the proper grammar.

References

[Amershi et al., 2014] Amershi, S., Cakmak, M., Knox, W. B., and Kulesza, T. (2014). Power to the people: The role of humans in interactive machine learning. *Ai Magazine*, 35(4):105–120. 5

[Amershi et al., 2019] Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., et al. (2019). Guidelines for human-ai interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*, pages 1–13. 5

[Arnórsson et al., 2023] Arnórsson, S., Abeillon, F., Al-Hazwani, I., Bernard, J., Hauptmann, H., and El-Assady, M. (2023). Why am i reading this? explaining personalized news recommender systems. 6

[Bernard et al., 2015a] Bernard, J., Daberkow, D., Fellner, D., Fischer, K., Koepler, O., Kohlhammer, J., Runnwerth, M., Ruppert, T., Schreck, T., and Sens, I. (2015a). Visinfo: a digital library system for time series research data based on exploratory search—a user-centered design approach. *International Journal on Digital Libraries*, 16:37–59. 5

[Bernard et al., 2015b] Bernard, J., Sessler, D., Bannach, A., May, T., and Kohlhammer, J. (2015b). A visual active learning system for the assessment of patient well-being in prostate cancer research. In *Proceedings of the 2015 Workshop on Visual Analytics in Healthcare*, pages 1–8. 4

[Bernard et al., 2018] Bernard, J., Zeppelzauer, M., Sedlmair, M., and Aigner, W. (2018). Vial: a unified process for visual interactive labeling. *The Visual Computer (TVCJ)*, 34(9):1189–1207. 5

[Bradley et al., 2018] Bradley, A. J., El-Assady, M., Coles, K., Alexander, E., Chen, M., Collins, C., Jänicke, S., and Wrisley, D. J. (2018). Visualization and the digital humanities. *IEEE computer graphics and applications*, 38(6):26–38. 6

[Brown et al., 2012] Brown, E. T., Liu, J., Brodley, C. E., and Chang, R. (2012). Dis-function: Learning distance functions interactively. In *2012 IEEE conference on visual analytics science and technology (VAST)*, pages 83–92. IEEE. 5

[Cabrera et al., 2019] Cabrera, Á. A., Epperson, W., Hohman, F., Kahng, M., Morgenstern, J., and 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. IEEE. 5

[Endert et al., 2012] Endert, A., Fiaux, P., and North, C. (2012). Semantic interaction for visual text analytics. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 473–482. 5

[Endert et al., 2014] Endert, A., Hossain, M. S., Ramakrishnan, N., North, C., Fiaux, P., and Andrews, C. (2014). The human is the loop: new directions for visual analytics. *Journal of intelligent information systems*, 43:411–435. 2

[Gratzl et al., 2013] Gratzl, S., Lex, A., Gehlenborg, N., Pfister, H., and Streit, M. (2013). Lineup: Visual analysis of multi-attribute rankings. *IEEE transactions on visualization and computer graphics*, 19(12):2277–2286. 5

[Höferlin et al., 2012] Höferlin, B., Netzel, R., Höferlin, M., Weiskopf, D., and Heidemann, G. (2012). Inter-active learning of ad-hoc classifiers for video visual analytics. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 23–32. IEEE. 5

[Jeong et al., 2009] Jeong, D. H., Ziemkiewicz, C., Fisher, B., Ribarsky, W., and Chang, R. (2009). ipca: An interactive system for pca-based visual analytics. In *Computer Graphics Forum*, volume 28, pages 767–774. Wiley Online Library. 5

[Kamalzadeh et al., 2016] Kamalzadeh, M., Kralj, C., Möller, T., and Sedlmair, M. (2016). Tagflip: active mobile music discovery with social tags. In *Proceedings of the 21st international conference on intelligent user interfaces*, pages 19–30. 5

[Kappas, 2013] Kappas, A. (2013). Social regulation of emotion: Messy layers. *Frontiers in Psychology 1664-1078*, 4. 3

[Lan et al., 2023] Lan, X., Wu, Y., and Cao, N. (2023). Affective visualization design: Leveraging the emotional impact of data. *IEEE Transactions on Visualization and Computer Graphics*. 4, 6

[Liao and Varshney, 2021] Liao, Q. V. and Varshney, K. R. (2021). Human-centered explainable ai (xai): From algorithms to user experiences. *arXiv preprint arXiv:2110.10790*. 5

[Lupi, 2017] Lupi, G. (2017). Data humanism: the revolutionary future of data visualization. *Print Magazine*, 30(3). 1, 2, 6

[Mamani et al., 2013] Mamani, G. M., Fatore, F. M., Nonato, L. G., and Paulovich, F. V. (2013). User-driven feature space transformation. In *Computer Graphics Forum*, volume 32, pages 291–299. Wiley Online Library. 5

[Pandey et al., 2023] Pandey, S., McKinley, O. G., Crouser, R. J., and Ottley, A. (2023). Do you trust what you see? toward a multidimensional measure of trust in visualization. In *2023 IEEE Visualization and Visual Analytics (VIS)*, pages 26–30. IEEE. 5, 6

[Radensky et al., 2022] Radensky, M., Downey, D., Lo, K., Popovic, Z., and Weld, D. S. (2022). Exploring the role of local and global explanations in recommender systems. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–7. 4

[Ritter et al., 2018] Ritter, C., Altenhofen, C., Zeppelzauer, M., Kuijper, A., Schreck, T., and Bernard, J. (2018). Personalized visual-interactive music classification. In *EuroVA@ EuroVis*, pages 31–35. 5

[Sacha et al., 2017] Sacha, D., Sedlmair, M., Zhang, L., Lee, J. A., Peltonen, J., Weiskopf, D., North, S. C., and Keim, D. A. (2017). What you see is what you can change: Human-centered machine learning by interactive visualization. *Neurocomputing*, 268:164–175. 5

[Sachdeva et al., 2023] Sachdeva, M., Burmeister, J., Kohlhammer, J., and Bernard, J. (2023). Lfpeers: Temporal similarity search and result exploration. *Computers & Graphics*. 5

[Shneiderman, 2020] Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction*, 36(6):495–504. 4

[So et al., 2020] So, W., Bogucka, E. P., Šćepanović, S., Joglekar, S., Zhou, K., and Quercia, D. (2020). Humane visual ai: Telling the stories behind a medical condition. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):678–688. 6

[Tintarev and Masthoff, 2007] Tintarev, N. and Masthoff, J. (2007). Effective explanations of recommendations: user-centered design. In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 153–156. 5

[Wang et al., 2022] Wang, Z. J., Kale, A., Nori, H., Stella, P., Nunnally, M. E., Chau, D. H., Vorvoreanu, M., Wortman Vaughan, J., and Caruana, R. (2022). Interpretability, then what? editing machine learning models to reflect human knowledge and values. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4132–4142. 2

[Wickham, 2010] Wickham, H. (2010). A layered grammar of graphics. *Journal of Computational and Graphical Statistics*, 19(1):3–28. 2