MLVis: Machine Learning Methods in Visualisation for Big Data (2024)
Daniel Archambault, Ian Nabney, and Jaakko Peltonen (Editors)

User-Adaptive Visualizations: An Exploration with GPT-4

F. Yanez and C. Nobre

University of Toronto, Canada

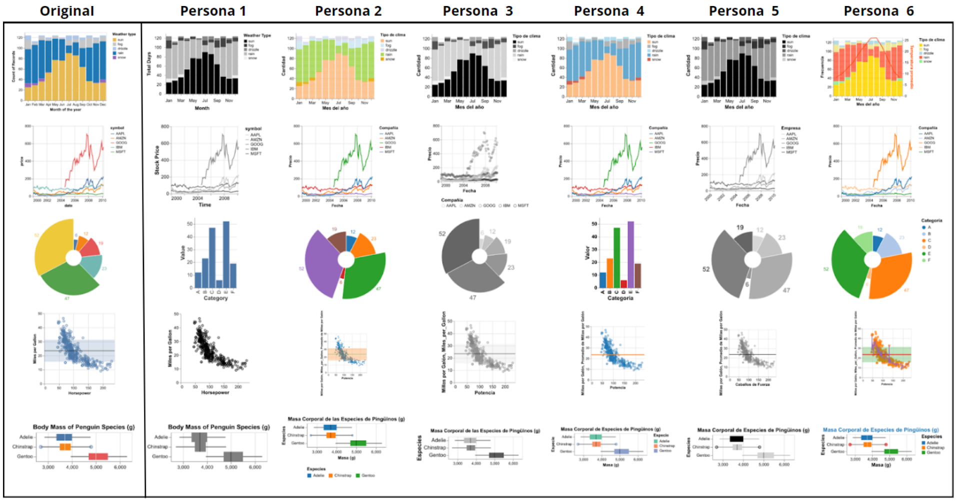

Figure 1: *Comparative analysis of data visualizations tailored to artificial user personas. The first column showcases the original visualization set, while subsequent columns display customized versions for Personas 1 through 6.*

Abstract

Data visualizations aim to enhance cognition and data interpretation. However, individual differences impact visual analysis, suggesting a personalized approach may be more effective. Current efforts focus on the study of generating visualizations with Large Language Models, lacking the user personalization component. This project explores using such models, specifically GPT-4, for modifying data visualizations to tailor to individual user characteristics. We developed a study to test GPT-4's ability to generate personalized visualizations. Statistical analysis of our results shows that for some personas, GPT is effective at personalizing the visualization. However, not all personalizations led to statistically significant improvements, suggesting variability in the effectiveness of LLM-driven personalization. These findings underline the importance of further exploring how personalized visualizations can best meet diverse user needs.

CCS Concepts

• ***Human-centered computing*** → ***Visualization design and evaluation methods;*** *Information visualization; User models;*

1. Introduction

One of the main goals of data visualizations is to enhance human cognition and assist in processing large amounts of data. However, recent studies have revealed that the process of visual analysis differs between individuals [Ott20]. Individual differences in cognitive abilities, personality traits, and experiences shape how one perceives and processes visual information, often measured by outcomes like users' accuracy, speed, and attention when performing tasks (e.g., [CCTL15,COSM20,OCZC15]). This indicates that personalizing visualizations to meet specific needs and preferences of users can be advantageous given that a one-size-fits-all approach may not be optimal.

The field of *user-adaptive visualizations* studies how visualizations can adjust to the characteristics and inclinations of individual users [CCTL15]. Prior work provides evidence on how adaptation can enhance the overall effectiveness of the visual analysis process (e.g., [LWC20]), reduce cognitive overload [ZK09, PYO*13], and improve data discovery rates [MHN*22].

Acknowledging the opportunity to enhance the visual analysis process, we explore the capacity of a Large Language Model (LLM), specifically GPT-4 [AAA*23], to generate visualizations that go beyond general charts by making modifications tailored to the individual's unique lens.

DOI: 10.2312/mlvis.20241126

Our study aims to harness the power of GPT-4 to modify data visualizations customized to user traits. Specifically, our research question explores whether GPT-4 can successfully create visualizations that are tailored to the unique combinations of individual characteristics of different user profiles. We also investigate how these AI-generated visualizations compare to original untailored ones in terms of aligning with the needs of diverse user personas. Thus, this work's contributions are twofold: firstly, it demonstrates the practicality of employing GPT-4 for the personalized generation of visualizations. Secondly, it offers valuable insights into the limitations and capabilities of GPT-4 in adapting to specific user requirements. Although this proof of concept does not involve real-world users, it lays the groundwork for subsequent empirical investigations that can validate the effectiveness of personalized visualizations in real-world scenarios.

2. Related Work

The exploration of LLMs in the domain of data visualization represents a burgeoning field of study, merging the intricacies of human-centric design with the computational prowess of artificial intelligence. In this section, we explore previous work on the effectiveness of GPT-4 for generating visualizations from language queries, performing visual analysis tasks, and creating grammar-agnostic visualization.

Generating visualizations from natural language was investigated by Maddigan & Susnjak [MS23], with Chat2Vis, a system that converts natural language queries into data visualizations. It leverages the capabilities of LLMs such as ChatGPT and GPT-3 to address the challenges of ambiguous natural language, improving the reliability and accuracy of visualization generation. This approach contrasts with traditional methods, offering a cost-effective and efficient solution by reducing the need for tailored grammar rules and specialized NLP tools. Lei Wang et al. [WZW*23] introduced LLM4Vis, a system for visualization recommendation that leverages ChatGPT's capabilities to provide natural and explainable recommendations. Unlike traditional machine learning approaches that require extensive training datasets, LLM4Vis operates effectively in few-shot and zero-shot settings by using a small number of demonstration examples to guide the model. Later, Li et al. (2024) [LWA*24], demonstrated the effectiveness of GPT-3.5 in interpreting natural language queries to generate precise visualization specifications. Employing Vega-Lite grammar and a comprehensive evaluation framework, their research not only underscores the advanced capabilities of LLMs in surpassing traditional methods but also highlights the nuanced improvements achieved through few-shot learning approaches.

Parallel to these advancements, Vázquez's (2024) [Váz24] research extends the application of LLMs to a broader spectrum of visualization tasks, probing the models' versatility across different libraries and their adaptability in customizing visual representations. This work illuminates the inherent potential of LLMs to cater to diverse visualization demands, yet it also surfaces critical challenges in model flexibility and the depth of customization. Such insights are invaluable in understanding the landscape of LLM-enabled visualization generation and setting the stage for further innovation.

In the context of tailoring visualizations LIDA [Dib23] provides a pivotal shift towards automating the visualization generation process by innovating the use of LLMs alongside Image Generation Models to automate the creation of both grammar-agnostic visualizations and intricate infographics. This tool not only streamlines the visualization process for novices but also sets a new benchmark in the field by demonstrating the potential of AI in intuitively navigating the complexities of data representation.

Our work builds on existing work by studying how visualizations can be personalized to align with specific user characteristics. By focusing on the interplay between user traits and visualization efficacy, it aims to extend the utility of LLMs in creating user-tailored visual representations of data. In doing so, it contributes a unique perspective to the ongoing discourse, underscoring the importance of user-centric approaches in the evolving narrative of data visualization.

3. Study Design

Our study employs a two-part approach designed to scrutinize the role of GPT-4 in generating user-centric visualizations. We first simulate artificial user profiles, which involve crafting a series of diverse, hypothetical user personas. These personas are equipped with varied cognitive abilities, personality traits, and preferences to mimic a broad user spectrum. We then provide GPT-4 with a curated set of visualizations which GPT-4 is tasked with adapting to the user personas. Here, the model's knowledge of visualization literature is put to the test, as it attempts to produce outputs that align with the fabricated personas' characteristics. In the sections below, we describe the specific steps and parameters used in the study.

Selection of Visualizations. We selected an initial set of 25 visualizations from the Vega-Lite example gallery[†], that ensured a wide variety of chart types and complexities. Following initial test iterations with GPT-4 to adjust visualizations based on user characteristics, we excluded those prone to rendering errors. From the remaining visualization, we selected 5 that represented common, yet varied types of charts. They are: stacked bar chart, multi-series line chart, radial plot, scatterplot with mean and standard deviation overlay, box plot with pre-calculated summaries (see Figure 1).

Simulation of Artificial User Profiles. For this study, we simulated a set of artificial user profiles, which the visualizations would adapt to. We created these profiles intending to mirror the rich diversity of human traits and preferences in real-world scenarios.

Defining User Characteristics: drawing from previous work on user-adaptive visualization literature, we chose to include in this study a combination of the most studied long-term and short-term characteristics. These characteristics have been identified as having the highest impact on research in user-adaptive visualizations, making them essential for developing comprehensive adaptation strategies. Our selection is grounded in an extensive literature review, focusing on characteristics that have been the most relevant to research in user-adaptive visualizations. Table 1 contains a list of these characteristics, each accompanied by multiple example references where they have been studied in various contexts.

† https://vega.github.io/vega-lite/examples/

Term	Category	Subcategory	Values defined for this study
Short	Cognitive States	Attention [WLMB*14, AAGP23]	Focused, Divided
		Confusion [AHM16, CAGM22]	High, Low
		Cognitive Load [AHM16, BWH*23]	Overloaded, Optimal
Long	Personality Traits	Locus of Control [OCZC15, ARG*20, SLC*20, DAG22]	Internal, External
		Need for Cognition [SLC*20, LTC21]	Low, High
		Neuroticism and Extraversion (N&E) [ARG*20, DAG22]	Introverted, Extraverted
	Cognitive Abilities	Perceptual Speed [TCCH12, CLRT17, CLRT20, AAGP23]	Slow, Fast
		Visual Working Memory (WM) [TCCH12, CLRT17, CLRT20, AAGP23]	Low Capacity, High Capacity
		Spatial Ability [MDF12, CLRT20]	Low, High
	Experience/ Bias	Visualization Literacy [CLRT17, CLRT20]	Low, High
		Data Expertise [TCCH12, AAGP23]	Novice, Expert
	Visual Preferences	Color Scheme	Bright Colors, Monochrome
		Visual Complexity	Simple, Complex
	Demographics	Age	Youth, Adult, Senior
		Color Perception	Normal Eyesight, Color Blind
		Language	English, Spanish

Table 1: *User characteristics and specific values included in this study, based on previous works related to data visualization.*

Generating User Profiles: with the aforementioned user characteristics, we meticulously synthesized ten multi-dimensional personas as shown in Table 2. These profiles were algorithmically created with uniform randomness from the values per subcategory defined for this study, ensuring diverse profiles.

We closely reviewed the ten artificial personas and selected six that were significantly different from each other, eliminating any that were too similar or had inconsistent traits. This careful selection aimed to cover a wide range of potential users, providing a solid foundation for our tests an LLM to generate custom visualizations.

Visualization Adaptation Process. This process involved two key steps: (1) providing tailored prompts to GPT-4 and (2) the autonomous adaptations of visualizations by the model. Each chart's modification was guided by a comprehensive prompt, detailing the user's specific characteristics. The GPT-4 instance, leveraging its vast knowledge base on information visualization literature, decided on both aesthetic and data-driven changes that aligned with user preferences.

The process was powered by OpenAI's Chat Completion API for the *gpt-4-turbo-preview* model (parameters used were a temperature of 0.2 and top_p equal to 1), calibrated to generate visualizations using Vega-Lite specifications. The model was instantiated with the following system prompt: *"You are a helpful assistant and an expert data analyst capable of creating or modifying data visualizations in the form of Vega-Lite specs. The modifications can be on aesthetics as well as the information shown. Keep in mind known best practices in information visualization."* The model operated autonomously within the defined system prompt constraints, focusing on both aesthetics and data representation with a grounding in best practices. This procedure was serial, treating each visualization instance as a unique task without leveraging historical data, ensuring that each outcome was independently tailored to the respective persona.

4. Experimentation and Evaluation

For the experimental setup, we utilized the OpenAI's Chat Completion API, specifically the *"gpt-4-vision-preview"* model, with its default parameters to evaluate both original and tailored visualizations, following a self-evaluation approach similar to prior work (e.g. [Dib23]). This process involved creating ten different instances of the model for each set of user characteristics—ensuring that the Central Limit Theorem for the mean of their scores, over the five different visualizations, can be applied—, and assessing how well each visualization adapted to the given persona. This methodology, akin to a between-subjects study, allowed for a controlled comparison between original and tailored visualizations, ensuring robustness through the replication of artificial personas to accommodate the stochastic nature of LLM outcomes.

In our evaluation technique, we leveraged GPT-4 instances to assess how well each visualization was personalized according to the predefined user characteristics. Each persona was presented with a visualization and asked to rate its adaptation on a scale from 1 (minimal adaptation) to 100 (optimal adaptation). The results included a score and its rationale. The score enabled a direct evaluation of how effectively each visualization aligned with the hypothetical user profiles, providing a clear metric for gauging personalization success without the need for additional subjective interpretation. The comparison was rigorously conducted using *scipy.stats*' T-test for the means of two independent samples of scores for each persona, as well as a comprehensive test across every persona, to identify any significant effects attributable to the personalization process.

5. Results and Discussions

Statistical Analysis: The statistical analysis (T-test for the means of two independent samples of scores for each persona with an assumption of normality with equal variances) reveals a nuanced effect of personalization on visualization effectiveness as shown in Table 3. For instance, the significant improvement in scores for Personas 1, 5, and 6 (p-values < 0.05) underscores the GPT-4's ability to significantly enhance visualizations for users with different characteristics. However, the results for Personas 2, 3, and 4, with p-values well above the significance threshold, suggest that personalization might not always lead to a statistically significant improvement. This outcome, however, could be related to the modified charts being quite similar to the original one based on the user traits. This mixed outcome invites further investigation into the

Subcategory	**Persona 1**	**Persona 2**	**Persona 3**	**Persona 4**	**Persona 5**	**Persona 6**
Attention	Divided	Divided	Focused	Focused	Divided	Divided
Confusion	Low	Low	Low	High	Low	High
Cognitive Load	Overloaded	Overloaded	Optimal	Overloaded	Overloaded	Optimal
Locus of Control	Internal	Internal	Internal	Internal	External	External
Need for Cognition	High	High	Low	High	Low	High
N&E	Introverted	Extraverted	Introverted	Extraverted	Introverted	Extraverted
Perceptual Speed	Fast	Fast	Slow	Slow	Slow	Fast
Visual WM	Low Capacity	High Capacity	Low Capacity	Low Capacity	High Capacity	High Capacity
Spatial Ability	Low	Low	Low	High	Low	Low
Visualization Literacy	Low	High	Low	Low	Low	High
Data Expertise	Novice	Novice	Novice	Novice	Novice	Expert
Color Scheme	Monochrome	Bright Colors	Monochrome	Bright Colors	Monochrome	Bright Colors
Visual Complexity	Simple	Complex	Complex	Simple	Complex	Complex
Age	Senior	Adult	Senior	Youth	Senior	Senior
Color Perception	Color Blind	Color Blind	Normal Eyesight	Color Blind	Normal Eyesight	Normal Eyesight
Language	English	Spanish	Spanish	Spanish	Spanish	Spanish

Table 2: *Detailed breakdown of six artificially created personas showcasing a diverse range of user characteristics.*

conditions under which LLM-driven personalization most effectively meets user needs, potentially guiding future enhancements in adaptive visualization technologies.

Discussion on Personalization: Our observations reveal the LLM's nuanced approach to personalizing visualizations, as shown in Figure 1. For instance, for Personas 1 and 4, with low data and visualization literacy as well as high cognitive load (Table 2), GPT-4 simplified complex charts, such as transforming radial charts into bar charts and removing extraneous details from the scatter and box plots. Conversely, for Persona 6, with optimal cognitive capacity and high literacy, the model enriched the visualizations with brighter colors and additional data elements like trend lines and legends. These adjustments underscore the LLM's capacity to tailor visual representations to accommodate the specific cognitive and aesthetic preferences of diverse personas, showcasing a significant stride towards more accessible and engaging data visualizations.

Discussion on Chart Complexity: Along with the score for each visualization, we asked the artificial user to provide a rationale for its answer. Despite instructions not to consider chart type preferences, GPT-4 often factored this into its scoring. Thus, traits with a preference for high visual complexity or a high need for cognition scored complex charts higher. At the same time, users with the opposite traits would argue that the modification was not as effective because the chart type was still too complex for them.

ID	μ_0	s_0	μ_1	s_1	t	p-value
1	41.40	14.81	**52.24**	14.90	-3.65	**2.12e-04**
2	**69.06**	7.42	67.72	7.54	0.90	8.14e-01
3	**41.82**	14.72	41.10	8.53	0.30	6.17e-01
4	43.90	12.01	**47.10**	11.13	-1.38	8.50e-02
5	39.20	10.27	**42.80**	9.21	-1.85	**3.40e-02**
6	74.20	5.24	**76.66**	5.79	-2.23	**1.41e-02**
Overall	51.60	18.20	**54.60**	16.48	-2.23	**1.71e-02**

Table 3: *Statistical overview highlighting the impact of LLM-personalized visualizations on artificial-users scores, indicating the significance of differences between original and tailored ones.*

6. Limitations and Future Work

We recognize that the study design (i.e., focused on utilizing a fixed set of user characteristics to simulate personas) inherently limits its ability to grasp the full breadth of individual user preferences and nuances. This constraint is further compounded by the artificial nature of these personas, which, while instrumental for controlled experimentation, may not authentically replicate the complexity of real-world user interactions. Similarly, this study does not consider the specific analysis tasks that users would need to perform with the visualizations, which can significantly impact the effectiveness of the visual designs. In this sense, the user's expertise level, particularly related to the task at hand, could affect the scoring process. Moreover, while numerous user characteristics might offer valuable insights into user-adaptive visualization (e.g., visual acuity, perception of contrast), it is not feasible to incorporate every potentially relevant characteristic into a single study due to constraints on scope and resources. However, expanding the scope of user characteristics beyond the six personas could provide a more detailed and encompassing understanding of diverse user needs.

Addressing these limitations by broadening the variety of chart types and pre-training or fine-tuning the model with a deeper focus on data visualization best practices and user characteristics' impact on interaction could substantially improve the project's applicability and impact. At the same time, future studies could aim to incorporate practical usage scenarios and user expertise to better assess how personalized visualizations aid in real-world data analysis tasks.

7. Conclusion

This research explores the domain of personalized data visualization using GPT-4, demonstrating the potential of LLMs to tailor visualizations to user characteristics. Through the creation and evaluation of artificial personas, this study highlighted both the capabilities and limitations of current models in generating user-centric visual designs. Looking forward, integrating real user characteristics and enhancing model training presents a promising path to more effective personalized visualizations. This work lays a foundation for future exploration, aiming to further bridge the gap between user-centered tailoring and data visualization.

References

[AAA*23] ACHIAM J., ADLER S., AGARWAL S., AHMAD L., AKKAYA I., ALEMAN F. L., ALMEIDA D., ALTENSCHMIDT J., ALTMAN S., ANADKAT S., ET AL.: Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023). 1

[AAGP23] ANDREOU P., AMYROTOS C., GERMANAKOS P., POLYCARPOU I.: Human-centered Information Visualization Adaptation Engine. In *Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization* (2023), UMAP '23, Association for Computing Machinery, pp. 25–33. URL: https://dl.acm.org/doi/10.1145/3565472.3592961, doi:10.1145/3565472.3592961. 3

[AHM16] ASENOV D., HILLIGES O., MÜLLER P.: The Effect of Richer Visualizations on Code Comprehension. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), CHI '16, Association for Computing Machinery, pp. 5040–5045. URL: https://dl.acm.org/doi/10.1145/2858036.2858372, doi:10.1145/2858036.2858372. 3

[ARG*20] ALVES T., RAMALHO B., GONÇALVES D., GAMA S., HENRIQUES-CALADO J.: Exploring How Personality Models Information Visualization Preferences. In *2020 IEEE Visualization Conference (VIS)* (2020), pp. 201–205. doi:10.1109/VIS47514.2020.00047. 3

[BWH*23] BANCILHON M., WRIGHT A., HA S., CROUSER R. J., OTTLEY A.: Why Combining Text and Visualization Could Improve Bayesian Reasoning: A Cognitive Load Perspective. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023), CHI '23, Association for Computing Machinery, pp. 1–. URL: https://dl.acm.org/doi/10.1145/3544548.3581218, doi:10.1145/3544548.3581218. 3

[CAGM22] CENEDA D., ARLEO A., GSCHWANDTNER T., MIKSCH S.: Show Me Your Face: Towards an Automated Method to Provide Timely Guidance in Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics 28*, 12 (2022), 4570–4581. doi:10.1109/TVCG.2021.3094870. 3

[CCTL15] CONATI C., CARENINI G., TOKER D., LALLÉ S.: Towards User-Adaptive Information Visualization. *Proceedings of the AAAI Conference on Artificial Intelligence 29*, 1 (2015). URL: https://ojs.aaai.org/index.php/AAAI/article/view/9775, doi:10.1609/aaai.v29i1.9775. 1

[CLRT17] CONATI C., LALLÉ S., RAHMAN M. A., TOKER D.: Further Results on Predicting Cognitive Abilities for Adaptive Visualizations. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence* (2017), International Joint Conferences on Artificial Intelligence Organization, pp. 1568–1574. URL: https://www.ijcai.org/proceedings/2017/217, doi:10.24963/ijcai.2017/217. 3

[CLRT20] CONATI C., LALLÉ S., RAHMAN M. A., TOKER D.: Comparing and Combining Interaction Data and Eye-tracking Data for the Real-time Prediction of User Cognitive Abilities in Visualization Tasks. *ACM Transactions on Interactive Intelligent Systems 10*, 2 (2020), 1–41. URL: https://dl.acm.org/doi/10.1145/3301400, doi:10.1145/3301400. 3

[COSM20] CROUSER R. J., OTTLEY A., SWANSON K., MONTOLY A.: Investigating the role of locus of control in moderating complex analytic workflows. *EuroVis 2020-Short Papers* (2020). 1

[DAG22] DELGADO T., ALVES T., GAMA S.: How neuroticism and locus of control affect user performance in high-dimensional data visualization. *Computers & Graphics 109* (2022), 88–99. URL: https://www.sciencedirect.com/science/article/pii/S0097849322001844, doi:10.1016/j.cag.2022.10.006. 3

[Dib23] DIBIA V.: LIDA: A Tool for Automatic Generation of Grammar-Agnostic Visualizations and Infographics using Large Language Models, 2023. URL: http://arxiv.org/abs/2303.02927, arXiv:2303.02927. 2, 3

[LTC21] LALLÉ S., TOKER D., CONATI C.: Gaze-Driven Adaptive Interventions for Magazine-Style Narrative Visualizations. *IEEE Transactions on Visualization and Computer Graphics 27*, 6 (2021), 2941–2952. doi:10.1109/TVCG.2019.2958540. 3

[LWA*24] LI G., WANG X., AODENG G., ZHENG S., ZHANG Y., OU C., WANG S., LIU C. H.: Visualization generation with large language models: An evaluation. *arXiv preprint arXiv:2401.11255* (2024). 2

[LWC20] LALLÉ S., WU T., CONATI C.: Gaze-Driven Links for Magazine Style Narrative Visualizations. In *2020 IEEE Visualization Conference (VIS)* (2020), pp. 166–170. doi:10.1109/VIS47514.2020.00040. 1

[MDF12] MICALLEF L., DRAGICEVIC P., FEKETE J.-D.: Assessing the Effect of Visualizations on Bayesian Reasoning through Crowdsourcing. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2536–2545. doi:10.1109/TVCG.2012.199. 3

[MHN*22] MONADJEMI S., HA S., NGUYEN Q., CHAI H., GARNETT R., OTTLEY A.: Guided data discovery in interactive visualizations via active search. In *2022 IEEE Visualization and Visual Analytics (VIS)* (2022), IEEE, pp. 70–74. 1

[MS23] MADDIGAN P., SUSNJAK T.: Chat2vis: Generating data visualizations via natural language using chatgpt, codex and gpt-3 large language models. *IEEE Access 11* (2023), 45181–45193. doi:10.1109/ACCESS.2023.3274199. 2

[OCZC15] OTTLEY A., CROUSER R. J., ZIEMKIEWICZ C., CHANG R.: Manipulating and controlling for personality effects on visualization tasks. *Information Visualization 14*, 3 (2015), 223–233. URL: https://doi.org/10.1177/1473871613513227, doi:10.1177/1473871613513227. 1, 3

[Ott20] OTTLEY A.: *Adaptive and Personalized Visualization*. Synthesis Lectures on Visualization. Springer International Publishing, 2020. URL: https://link.springer.com/10.1007/978-3-031-02607-2, doi:10.1007/978-3-031-02607-2. 1

[PYO*13] PECK E. M. M., YUKSEL B. F., OTTLEY A., JACOB R. J., CHANG R.: Using fNIRS brain sensing to evaluate information visualization interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), ACM, pp. 473–482. URL: https://dl.acm.org/doi/10.1145/2470654.2470723, doi:10.1145/2470654.2470723. 1

[SLC*20] SHEIDIN J., LANIR J., CONATI C., TOKER D., KUFLIK T.: The effect of user characteristics in time series visualizations. In *Proceedings of the 25th International Conference on Intelligent User Interfaces* (2020), ACM, pp. 380–389. URL: https://dl.acm.org/doi/10.1145/3377325.3377502, doi:10.1145/3377325.3377502. 3

[TCCH12] TOKER D., CONATI C., CARENINI G., HARATY M.: Towards Adaptive Information Visualization: On the Influence of User Characteristics. In *User Modeling, Adaptation, and Personalization*, Masthoff J., Mobasher B., Desmarais M. C., Nkambou R., (Eds.), Lecture Notes in Computer Science. Springer, 2012, pp. 274–285. 3

[Váz24] VÁZQUEZ P.-P.: Are llms ready for visualization? *arXiv preprint arXiv:2403.06158* (2024). 2

[WLMB*14] WALDNER M., LE MUZIC M., BERNHARD M., PURGATHOFER W., VIOLA I.: Attractive Flicker — Guiding Attention in Dynamic Narrative Visualizations. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2456–2465. doi:10.1109/TVCG.2014.2346352. 3

[WZW*23] WANG L., ZHANG S., WANG Y., LIM E.-P., WANG Y.: Llm4vis: Explainable visualization recommendation using chatgpt. *arXiv preprint arXiv:2310.07652* (2023). 2

[ZK09] ZIEMKIEWICZ C., KOSARA R.: Preconceptions and Individual Differences in Understanding Visual Metaphors. *Computer Graphics Forum 28*, 3 (2009), 911–918. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2009.01442.x, doi:10.1111/j.1467-8659.2009.01442.x. 1