

From Data to Narrative: Visualizing Complex Phenomena through Human-AI Co-Creation

Chiara Ceccarini^{1,*}, Ami Liçaj², Elisa Matteucci² and Giovanni Delnevo¹

¹Department of Computer Science and Engineering, University of Bologna, Italy

²Department of Architecture, University of Florence, Italy

Abstract

In the face of global emergencies such as climate change, migration, and pollution, there is an urgent need for more effective tools to communicate complex issues and support collective understanding. This research explores the potential of Human-AI Collaborative (HAIC) Systems to address these challenges, emphasizing the synergy between human creativity, intuition, and social awareness, and the computational power and analytical capabilities of artificial intelligence. By combining these complementary strengths, the project aims to develop new models for translating complex data into intuitive and emotionally resonant visual representations. A key objective of this research is to design a tool that can be used not only by data visualization experts but also by individuals with scientific knowledge yet limited experience in visual communication. The system seeks to empower scientists, researchers, and activists to create clear, engaging, and emotionally impactful visualizations without requiring technical expertise in dataviz design. Through techniques such as algorithmic image segmentation and curated color palette application, the system generates unconventional, personalized visuals that enhance both engagement and comprehension. This approach demonstrates how human-machine collaboration can democratize access to knowledge, fostering more inclusive, perceptive, and impactful communication of urgent global phenomena.

Keywords

Data Visualization, Human-AI Collaborative Systems, Human-AI Interaction, GenAI, Co-Creation, Digital Sustainability

1. Introduction

Today's global landscape is shaped by urgent and multifaceted crises such as climate change, mass migration, and the accelerating degradation of natural ecosystems [1]. These crises are inherently interconnected, spanning across disciplines, borders, and temporal scales, making them difficult to grasp through conventional means of analysis or communication [2]. As scientific knowledge, environmental monitoring, and global reporting generate increasingly intricate datasets, the challenge lies not only in collecting and analyzing this information but in making it understandable, accessible, and actionable for diverse audiences, including policymakers, educators, and the general public [3].

In this context, the need for more effective tools to communicate the complexity and foster collective understanding has become increasingly critical [4]. Clear, engaging, and emotionally resonant communication can help bridge the gap between data and decision-making [5]. This is where data visualization emerges as a vital discipline. Traditionally used to support scientific inquiry and technical reporting, data visualization has evolved into a powerful medium for storytelling, advocacy, and public engagement. Transforming abstract numbers into visual forms enables people to perceive patterns, detect anomalies, and derive insights that would otherwise remain hidden in raw data [6].

More than ever, data visualization plays a strategic role not only within academic and professional circles but also in the broader sociopolitical arena [7]. It serves as a bridge between knowledge and perception, allowing individuals and communities to better understand the scope, urgency, and human impact of global phenomena. When thoughtfully designed, it can evoke empathy, prompt critical

HAIC 2025 - First International Workshop on Human-AI Collaborative Systems, October 25, 2025, Bologna, Italy

*Corresponding author.

✉ chiara.ceccarini6@unibo.it (C. Ceccarini); ami.licaj@unifi.it (A. Liçaj); elisa.matteucci@unifi.it (E. Matteucci); giovanni.delnevo@unibo.it (G. Delnevo)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

reflection, and inspire action—qualities that are essential in addressing the most pressing challenges of our time.

In this evolving landscape, unconventional data visualizations, those that depart from standard charts, graphs, and maps, are gaining traction as powerful tools for engagement and meaning-making [8]. These forms, which may include abstract compositions, hybrid artistic-scientific renderings, or visually expressive metaphors, offer alternative ways of seeing that can resonate more deeply with diverse audiences [9]. Unlike traditional data graphics, which prioritize accuracy and legibility, unconventional visualizations emphasize emotional impact, cultural relevance, and interpretative openness [10]. When applied thoughtfully, they can elicit curiosity, foster empathy, and provoke dialogue, particularly around complex, multidimensional issues such as sustainability [11], climate change [12] or migration [13], where emotional resonance and social context are integral to understanding. This expanded visual repertoire challenges normative assumptions about how data should be represented, opening new avenues for communication that are both affective and informative.

Despite its potential, designing effective data visualizations remains a complex and resource-intensive task. It requires a nuanced balance of aesthetics, clarity, and narrative cohesion—skills often beyond the reach of non-experts. In response to this gap, recent advances in Generative AI, particularly Large Language Models (LLMs), have begun to offer new possibilities. These models are capable of generating unconventional and emotionally compelling visual representations from textual input, lowering the barrier to entry for individuals lacking formal training in design or visual analytics. However, while LLMs can inspire creative and novel visual forms, they often lack the necessary semantic control and consistency, particularly in critical aspects such as color encoding. This limitation becomes especially problematic when color is not merely decorative but must convey precise, domain-specific information, such as temperature thresholds, pollutant concentrations, or population categories.

This research addresses this challenge by introducing a Human-AI Collaborative (HAIC) system that combines the generative capacities of LLMs with human oversight and semantic control. The project aims to develop models and applications that generate unconventional data visualizations by leveraging the complementary strengths of human intuition, aesthetic sensibility, and social awareness with the computational power and analytical precision of AI. A first step is the development of a web-based application that segments and overlays images using curated color palettes to transform abstract datasets into visually compelling and emotionally resonant representations. By structuring collaboration between humans and machines around their complementary strengths, human domain expertise, interpretative judgment, and AI's generative and computational capabilities, we aim to bridge the gap between expressiveness and precision in AI-assisted visual communication, offering a novel approach to data communication that prioritizes accessibility, personalization, and engagement.

The rest of the paper goes as follows. Section 2 contextualizes our research within existing literature. Then, Section 3 outlines the proposed approach. Section 4 presents the web application developed, while Section 5 presents and discusses the results of user testing. Finally, Section 6 summarizes our contributions and proposes directions for further research.

2. Related work

Recent years have seen a growing interest in systems that enable collaboration between humans and artificial intelligence across a wide range of domains [14]. These systems are designed not to replace human effort but to enhance it, supporting people in making better decisions, working more efficiently, and unlocking new creative possibilities. Most of the studies in this space focus on user-centered approaches, building and testing prototypes with real users to evaluate usability, performance, and engagement in realistic contexts [15, 16].

In the field of data analysis, for example, Xie et al. introduced HAICChart, a visualization system that blends AI-generated insights with user-driven exploration [17]. Rather than relying solely on automated outputs, users could actively switch between views suggested by the AI and those they created themselves. This fluid, back-and-forth interaction allowed for a more dynamic and engaging

experience. The system proved to be not only faster (1.8 times quicker than traditional tools) but also more effective, improving information recall by 21%. These results underscore how tightly integrated AI support can enrich exploratory data work without taking control away from the user.

Storytelling within data journalism presents another interesting use case. An example is Erato, a co-creative editing system that helps users build narratives from structured datasets [18]. Instead of generating entire stories automatically, Erato worked as a kind of creative writing partner: it provided blocks of text based on the data, which users could freely rearrange, edit, or rewrite. Participants reported that this process led to more coherent and engaging stories, while also reducing the time needed to produce them. The findings suggest that AI can be a valuable tool in narrative construction, especially when working within data-driven constraints.

Turning to the artistic domain, Kim et al. explored how AI can support creative expression through Colorbo, a collaborative system for coloring Mandala art [19]. The system proposed harmonious color palettes based on symmetry and aesthetic principles, while still giving users full freedom to choose and modify colors. Particularly for non-artists, Colorbo created a space where creativity felt accessible rather than intimidating. Users described the experience as enjoyable and even meditative, highlighting AI's potential not only as a productivity tool but also as a source of inspiration and emotional engagement.

A similar spirit of collaboration guided the development of DuetDraw [20]. This system responded in real time to users' sketches, suggesting complementary shapes and lines that could expand or enhance the drawing. Users described the interaction as playful and expressive, and particularly appreciated the system's contextual suggestions and ability to follow their creative lead. The tool encouraged exploration and experimentation and helped to break down the common fear of making mistakes when drawing.

AI collaboration also holds promise in storytelling aimed at younger audiences. StoryDrawer enabled children to co-create illustrated stories with the help of an AI partner [21]. The system could generate characters, complete scenes, and respond to the child's input with visual suggestions, stimulating creativity and engagement. Evaluations showed that children not only enjoyed the experience but also stayed focused longer and produced more imaginative narratives, pointing to exciting opportunities for AI in educational and developmental contexts.

In a more design-focused context, Antony and Huang introduced ID.8, a generative AI system for creating visual stories [22]. It supported users throughout the creative process: starting from a prompt or idea, the AI could generate illustrations, propose narrative directions, or help refine drafts. With a usability score of 77.25 on the System Usability Scale (SUS), users found the tool intuitive and helpful. It proved particularly valuable for overcoming creative blocks, making it a useful aid for both beginner storytellers and experienced visual designers alike.

Despite their differences in context and application, these systems share several common threads. Most were evaluated using mixed methods that combined measurable performance indicators with user feedback. Across the board, users reported increased efficiency, higher quality outcomes, whether in terms of accuracy, coherence, or aesthetics, and a general sense of satisfaction with the collaborative process.

Taken together, these studies reflect the growing maturity of human-AI collaboration, highlighting also a few challenges: some users experienced communication mismatches between their intentions and the system's responses, while others noted unpredictable behavior or limitations tied to specific domains. However, when designed thoughtfully, with attention to user needs, transparency, and interaction design, AI systems can act as effective partners rather than passive tools.

3. Methodology

The design of the proposed tool was informed by a preliminary focus group involving the authors, two designers, and two computer scientists. This interdisciplinary collaboration aimed to identify the key steps necessary to support users, particularly those with limited experience in data visualization, in creating effective and emotionally engaging visual representations of complex phenomena. The

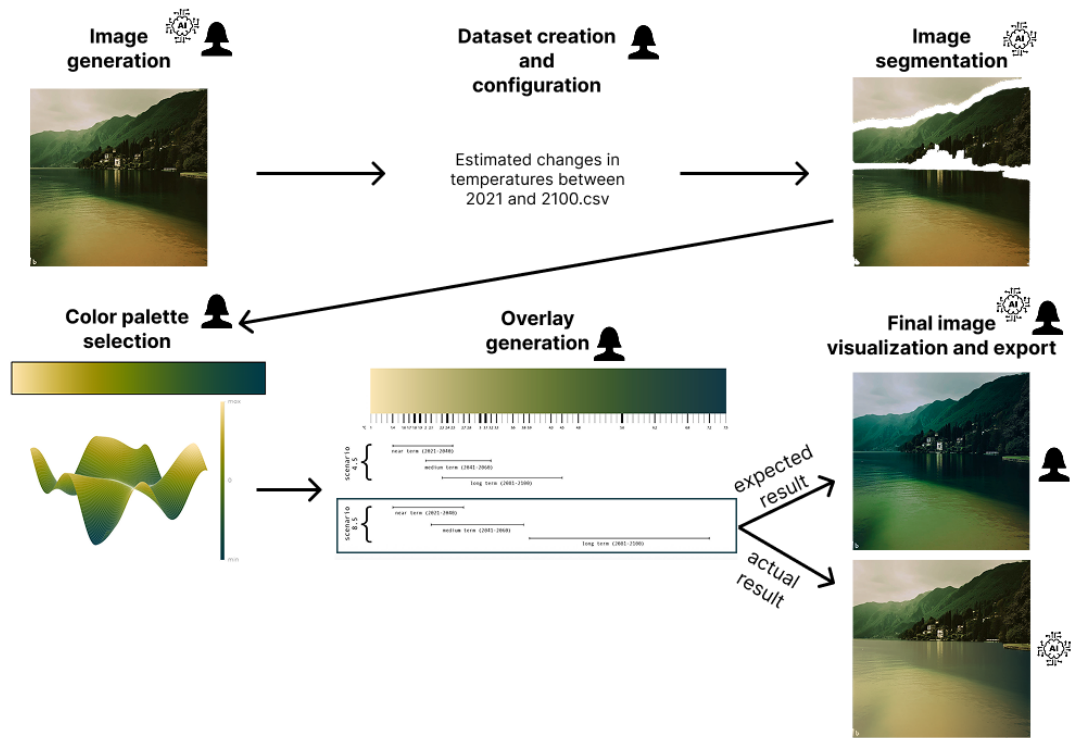


Figure 1: Our approach to creating data-driven visualizations leveraging a Human-AI co-creation tool. It consists of the following steps: (i) image generation, involving both human and AI; (ii) dataset creation and configuration, performed by the human; (iii) image segmentation, carried out by the AI; (iv) color palette selection, made by the human; (v) overlay generation, done by the human; and (vi) final image visualization and export, involving the human (top) and the AI (bottom).

discussion allowed the team to align on both functional requirements and user experience priorities, ensuring the resulting system would be intuitive, flexible, and supportive of non-expert users in the visualization process. Hence, to create data-driven visualizations to communicate complex phenomena, we propose an innovative approach based on human-AI co-creation. The approach, which relies on the strengths of both AI and humans, is composed of six steps, as visible in Figure 1. These steps structure the interaction between the user and the system, facilitating a guided yet creative workflow that enhances comprehension and emotional resonance through personalized, visually compelling outputs.

Image generation The first step of the process involves generating an image with the help of a generative AI system, aligned with the theme of the phenomenon to be analyzed and visualized (i.e., the raising of temperature or intensive farming), that will be the base for the unconventional visualization. This initial image serves as a conceptual and emotional anchor, translating abstract or complex issues into a visual form. Depending on the desired outcome and the specificity of the image, multiple interactions with the generative AI may be required, particularly if the user is not experienced in prompt engineering. This iterative dialogue between the user and the system allows for progressive refinement, ensuring the generated image effectively captures the intended meaning and emotional tone of the subject matter.

Dataset creation and configuration The second step consists of creating and configuring datasets. These datasets are the quantitative backbone representing the phenomenon of interest that we aim to convey to the users through our visualizations. They will serve as the quantitative basis for the visual overlays that will be applied to the images, transforming abstract information into perceivable insights.

Image segmentation Rather than applying a palette uniformly across the entire visual field, segmentation allows users to isolate specific areas of interest, such as an object or subject within a photograph, and enhance only those areas with data-driven visual elements. This approach not only improves readability and focus but also introduces a more personalized and meaningful method of interpreting data. After generating an image, users are prompted to select a region in which they wish to convey information using dataset values. This selection is made possible through semantic segmentation that automatically detects and delineates meaningful regions in the image (e.g., people, animals, landscapes), allowing users to select a specific segment with a single click. Once selected, the segmented area is treated as a “mask”, a binary or multi-class map that identifies which pixels should be subject to data-driven coloring and which should remain unaltered. This is crucial for maintaining the visual integrity of the original image while introducing interpretative layers where relevant.

Color palette selection After defining a dataset and segmenting an image, the user selects a color palette to represent the data. The idea is to use scientific palettes (e.g., those by Crameri [23]), as well as the option to create custom palettes by specifying gradient directions or start and end colors. This step ensures consistency and perceptual clarity in the visualization, as the selected palette maps data values to color in an intuitive and standardized way.

Overlay generation In this step, a gradient based on the selected dataset’s value range and the associated palette is generated. Crucially, this process offers the user significant control, providing the opportunity for manual mapping between specific dataset values and desired colors within the palette. This direct intervention allows users to finely tune how quantitative information is visually represented, ensuring that the chosen color transitions accurately reflect the nuances and significance of the underlying data according to domain-specific understanding. The gradient is computed and applied only to the masked region of the image. The overlay process leverages alpha blending techniques to integrate the gradient without disrupting the original visual context. The result is a composite image where the segmented area visually represents the underlying data in a non-conventional, yet informative way. This enhanced image provides users with a powerful tool for data interpretation and communication.

Final data-driven visualization and export The final step in the proposed methodology is the generation and export of the composite visualization. At this stage, the process culminates in the creation of the enhanced image, where the selected dataset’s values are visually integrated into the segmented region of the AI-generated base image. This resulting visualization represents the complete fusion of quantitative data with a compelling visual narrative. This last phase completes the data visualization workflow, enabling users to produce, evaluate, and share customized outputs tailored to specific analytical needs.

4. Web-based application

To operationalize the Human-AI co-creation methodology, we developed a modular, web-based application that guides users through each of the six core steps of the workflow described in Section 3. The system is structured around a sequential, screen-based interface, where each screen corresponds to a distinct stage of the visualization process. This includes image selection, dataset configuration, image segmentation, color palette definition, overlay generation, and final visualization export. The step-by-step design (see Figure 2) ensures an intuitive and accessible experience, particularly for users without prior expertise in data visualization or graphic design. By structuring the tool as a linear and intuitive journey, the application lowers the barrier to entry for users who may lack technical expertise in data visualization, while still offering flexibility and creative control. Each screen is tailored to support a specific phase of the process, ensuring clarity, consistency, and ease of use at every stage.

The application adopts a client-server architecture. The frontend is implemented using modern web technologies, primarily the Vue.js framework for reactive UI components and Bootstrap for responsive layout across devices. This architecture promotes usability, interactivity, and broad accessibility across platforms. The backend is developed in PHP for handling HTTP requests and interfacing with a MySQL database, which stores user-generated content, such as images, datasets, color palettes, and segmentation masks. The modular database schema enables scalability, supporting future extensions such as additional overlay types, palette libraries, or collaborative editing features.

The first step involves acquiring the base image for visualization. Users can leverage integrated OpenAI API access to generate images via iterative prompt refinement, aligning them with the phenomenon of interest. Alternatively, the system supports direct image upload, accommodating visuals sourced from other generative services or pre-existing assets. This dual approach ensures flexibility in obtaining the foundational visual for our unconventional visualizations.

The segmentation step is performed entirely on the client side using **DeepLab.js** [24], a JavaScript adaptation of Google’s DeepLab semantic segmentation model. Running on top of TensorFlow.js, this implementation enables in-browser segmentation without requiring any server-side image processing. Users can interactively select and isolate meaningful regions, such as people, objects, or landscapes, by clicking on detected segments. This not only enhances performance and privacy by avoiding the need to upload images to remote servers, but also enables real-time feedback and fine-tuned region selection.

To support precise data encoding, users can define both primary datasets and configurable sub-datasets. This functionality enables more granular mapping of numerical values onto image regions, enhancing flexibility in exploring and communicating multiple dimensions of the same phenomenon.

The overlay application leverages **OpenCV.js** [25], a WebAssembly-based port of the OpenCV computer vision library. OpenCV.js performs real-time gradient computation and pixel-level blending directly within the browser. Users can manually control the mapping of data values to color stops in the selected palette, ensuring that domain-specific semantics, such as threshold boundaries or critical values, are accurately represented. The final overlay is applied selectively to the segmented region using alpha blending techniques, preserving the integrity of the original image while introducing interpretative layers of quantitative meaning.

Overall, the application embodies the core principles of our methodology: guided co-creation, interpretability, and expressive visualization. By integrating state-of-the-art tools for segmentation and image processing in a browser-native environment, it offers a flexible yet powerful platform that lowers technical barriers while maintaining high visual and analytical fidelity. Its architecture is designed with extensibility in mind, paving the way for future integration of generative image creation, collaborative workflows, and richer forms of visual annotation.

5. Results and discussion

The aim of this study was to design and implement a tool capable of generating data-driven visualizations that could be used even by individuals without prior expertise in data visualization. The project followed a human-AI co-creation approach, allowing users, regardless of their technical background, to produce expressive, clear, and emotionally engaging visual narratives grounded in scientific data.

User testing played a central role in assessing the web tool’s usability and communicative effectiveness. The application was tested by five users, some of whom had little to no experience in data visualization. The tool was evaluated through a think-aloud protocol, in which participants were asked to verbalize their thoughts while interacting with the prototype. This approach provided valuable insights into user reasoning, usability challenges, and cognitive processes behind interface navigation, allowing for a comprehensive evaluation of the tool’s accessibility and its capacity to guide users through the visualization process. Overall, the results showed that the platform was largely perceived as intuitive, especially thanks to its step-by-step interface that broke the process into manageable stages.

The image segmentation phase consistently emerged as a particularly delicate and critical step, especially for non-expert users. A recurring challenge was their difficulty in discerning which segmented

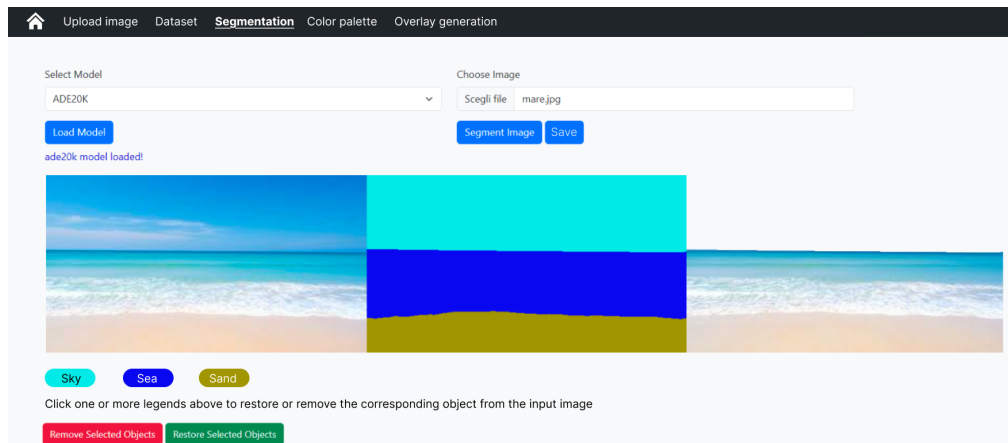


Figure 2: Segmentation screen of the web application. In this screen, users can apply semantic segmentation using a pre-trained model (e.g., ADE20K). The segmented image is displayed alongside the original, with each region labeled by color-coded categories (e.g., sky, sea, sand). Users can interactively select and remove or restore specific objects from the segmentation, tailoring the image to suit their visualization goals.

elements within an image held the most relevance for the specific dataset being visualized. During our testing, a significant issue was the varied precision of the segmentation output, particularly regarding the accurate recognition or distinction of finer details or less common objects in uploaded images (as depicted in Figure 1-image segmentation). This limitation is directly attributable to the specific pre-trained semantic segmentation model utilized. It is crucial to understand that segmentation models are trained on diverse and specialized datasets, consequently leading to distinct performance characteristics and domain-specificities. For instance, a model trained on ADE20K excels at broad scene understanding and object recognition across a wide array of categories, while one trained on Cityscapes is optimized for urban street scenes, and models trained on PASCAL VOC are highly proficient at segmenting common objects. Our current system's reliance on a general-purpose model, such as DeepLab trained solely on ADE20K, while robust, may not always align perfectly with the nuanced content or context of a user's image.

As a result, the segmentation sometimes introduced unintended or overly broad classifications, leading users to question the relevance or usefulness of certain regions within the visual. In some cases, users struggled to make meaningful associations between the segmented elements and the intended message, suggesting that this step still requires a degree of interpretation, contextual refinement, and domain awareness that current AI models alone cannot yet fully replicate.

Therefore, to overcome these limitations and enhance user experience, an advanced system should abstract this underlying model complexity. Instead of expecting users to understand the nuances of various training datasets, the system should intelligently suggest or automatically select the most appropriate segmentation model based on the inferred context of the user's image content and the phenomenon to be visualized. This context-aware model selection is paramount to achieving precise and meaningful segmentation, ensuring that the delineated regions accurately serve the data's intended narrative.

Additionally, the customization of the final visualization, particularly the choice of color palettes, overlays, and aesthetic parameters, sometimes proved difficult for users with no previous knowledge of visual storytelling or design principles. This highlights that, at least in the current version of the tool, human input remains essential to define the expressive tone and narrative impact of the output. As shown in Figure 1-final image, the image generated by the system (bottom image) appears very different from the one created by a designer to communicate the same data, resulting in a different emotional impact on the audience.

Nonetheless, users were able to understand the purpose of the tool and expressed appreciation for the creative and flexible experience it provided. Many recognized the potential of such a system for educational, scientific, or advocacy purposes. The combination of automation and user-driven control

was perceived as a core strength, fostering both accessibility and depth of interpretation.

These findings are consistent with earlier research on non-conventional data visualization, where user involvement is often crucial not only for usability improvements but also for maximizing communicative effectiveness [26].

6. Conclusion and Future Work

This work introduced an innovative web-based application as a first step to create personalized, emotionally engaging data visualizations through a guided process of human-AI collaboration. The tool was conceived to enable individuals with scientific expertise, but limited experience in visual design, to communicate complex phenomena in accessible and impactful ways.

The results of the user study confirm the validity of the proposed approach. The tool was generally perceived as easy to use and rich in creative potential. At the same time, the findings revealed that certain steps of the process, such as image segmentation and the definition of final visualization parameters, still require human interpretation and a basic understanding of visual principles. In its current form, the tool significantly lowers the entry barrier but does not yet allow for complete autonomy in the creation of effective visualizations without some form of guidance or prior knowledge.

Looking ahead, several directions for future development emerge from this study. One key area is the refinement of the segmentation interface and the incorporation of AI-supported suggestions to help users identify meaningful image elements based on context. Another promising direction is the integration of dynamic recommendations, templates, and usage examples, which could guide users in making stylistic and narrative decisions more confidently. Additional enhancements may include expanding the library of color palettes and visual styles, improving accessibility features, and potentially incorporating natural language input to make the tool even more inclusive.

In addition, testing multimodal models, such as LLaVA, Chameleon, and similar architectures that integrate both text and images, could represent a valuable step forward, improving usability and enabling the introduction of new functionalities. Equally important is the expansion of the user base, which would allow for the collection of more robust and representative statistics on application usage, ultimately guiding design choices and ensuring scalability. Future studies with larger and more diverse participant groups will be essential to assess the robustness, adaptability, and generalizability of the tool. Through these improvements, the project ultimately aims to support a more democratic and inclusive approach to data storytelling, one that leverages the power of artificial intelligence while embracing human sensitivity, creativity, and contextual awareness.

Acknowledgments

The theoretical framework on the integration of artificial intelligence, data distribution, and color associations was developed by the team from the Department of Architecture at the University of Florence, as part of the broader research line EcoVisualization strategies, which includes several studies on the role of data visualization and the principles of sustainability. The technical development of the web application and the implementation of the computer science components were carried out by the group of the Department of Computer Science and Engineering at the University of Bologna.

We would also like to thank Marco Costantini for his support in the development of the application, as well as the users who participated in the testing phase.

Declaration on Generative AI

During the preparation of this work, the authors used GPT-4 in order to: grammar and spell check. Further, the author(s) used Bing Image generation for part of Figure 1 in order to generate the image of the lake. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] C. Fletcher, W. J. Ripple, T. Newsome, P. Barnard, K. Beamer, A. Behl, J. Bowen, M. Cooney, E. Crist, C. Field, K. Hiser, D. M. Karl, D. A. King, M. E. Mann, D. P. McGregor, C. Mora, N. Oreskes, M. Wilson, Earth at risk: An urgent call to end the age of destruction and forge a just and sustainable future, *PNAS Nexus* 3 (2024). URL: <http://dx.doi.org/10.1093/pnasnexus/pgae106>. doi:10.1093/pnasnexus/pgae106.
- [2] F. Mazzocchi, Tackling modern-day crises: Why understanding multilevel interconnectivity is vital, *BioEssays* 43 (2020). URL: <http://dx.doi.org/10.1002/bies.202000294>. doi:10.1002/bies.202000294.
- [3] J. P. Shonkoff, Making developmental science accessible, usable, and a catalyst for innovation, *Applied Developmental Science* 24 (2018) 37–42. URL: <http://dx.doi.org/10.1080/10888691.2017.1421430>. doi:10.1080/10888691.2017.1421430.
- [4] P. M. Regular, G. J. Robertson, R. Rogers, K. P. Lewis, Improving the communication and accessibility of stock assessment using interactive visualization tools, *Canadian Journal of Fisheries and Aquatic Sciences* 77 (2020) 1592–1600. URL: <http://dx.doi.org/10.1139/cjfas-2019-0424>. doi:10.1139/cjfas-2019-0424.
- [5] A. H. Toomey, Why facts don't change minds: Insights from cognitive science for the improved communication of conservation research, *Biological Conservation* 278 (2023) 109886. URL: <http://dx.doi.org/10.1016/j.biocon.2022.109886>. doi:10.1016/j.biocon.2022.109886.
- [6] P. Sarkar, Y. D. Priya, P. B. Patel, P. Chatterjee Biswas, S. Arigela, S. Sallaram, Data visualization in transforming raw data into compelling visual narratives, in: 2024 International Conference on Trends in Quantum Computing and Emerging Business Technologies, IEEE, 2024, p. 1–6. URL: <http://dx.doi.org/10.1109/TQCEBT59414.2024.10545256>. doi:10.1109/tqcebt59414.2024.10545256.
- [7] T. U. Nærland, The political significance of data visualization: Four key perspectives, Amsterdam University Press, 2020. URL: http://dx.doi.org/10.5117/9789463722902_CH04. doi:10.5117/9789463722902_ch04.
- [8] V. R. Lee, M. H. Wilkerson, Data use by middle and secondary students in the digital age: A status report and future prospects (2018).
- [9] G. Lupi, Data humanism: The revolutionary future of data visualization, <https://www.printmag.com/article/data-humanism-future-of-data-visualization/>, 2017. [Accessed 09-06-2025].
- [10] D. Yan, C. Tang, S. Chandrasegaran, G. Kortuem, Reciportrait: a data humanism approach for collaborative sensemaking of personal data, in: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, pp. 1–21.
- [11] C. Ceccarini, T. Zambon, N. De Luigi, C. Prandi, Sdgs like you have never seen before!: Co-designing data visualization tools with and for university students, in: Proceedings of the 2023 ACM Conference on Information Technology for Social Good, GoodIT '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 521–529. URL: <https://doi.org/10.1145/3582515.3609577>. doi:10.1145/3582515.3609577.
- [12] C. Ceccarini, M. Ferreira, C. Prandi, N. Nunes, V. Nisi, Unusual suspects-visualizing unusual relationships of complex social phenomena with climate change, in: Proceedings of the 2023 ACM Conference on Information Technology for Social Good, 2023, pp. 494–503.
- [13] R. Risam, Beyond the migrant “problem”: Visualizing global migration, *Television & New Media* 20 (2019) 566–580.
- [14] H. J. Wilson, P. R. Daugherty, Collaborative intelligence: Humans and ai are joining forces, *Harvard business review* 96 (2018) 114–123.
- [15] T. Jiang, Z. Sun, S. Fu, Y. Lv, Human-ai interaction research agenda: A user-centered perspective, *Data and Information Management* 8 (2024) 100078. URL: <http://dx.doi.org/10.1016/j.dim.2024.100078>. doi:10.1016/j.dim.2024.100078.
- [16] G. Tumedei, C. Ceccarini, I. C. Jimenez Navarro, C. Prandi, From drawings to awareness: Exploring narrative visualization and ai to teach children about the fragile ecosystem of the mar menor

- lagoon, in: Proceedings of the 2025 ACM Designing Interactive Systems Conference, DIS '25, Association for Computing Machinery, New York, NY, USA, 2025, p. 2684–2700. URL: <https://doi.org/10.1145/3715336.3735722>. doi:10.1145/3715336.3735722.
- [17] Y. Xie, Y. Luo, G. Li, N. Tang, Haichart: Human and ai paired visualization system, Proceedings of the VLDB Endowment 17 (2024) 3178–3191. URL: <http://dx.doi.org/10.14778/3681954.3681992>. doi:10.14778/3681954.3681992.
 - [18] M. Sun, L. Cai, W. Cui, Y. Wu, Y. Shi, N. Cao, Erato: Cooperative data story editing via fact interpolation, IEEE Transactions on Visualization and Computer Graphics (2022) 1–11. URL: <http://dx.doi.org/10.1109/TVCG.2022.3209428>. doi:10.1109/tvcg.2022.3209428.
 - [19] E. Kim, J. Hong, H. Lee, M. Ko, Colorbo: Envisioned mandala coloring through human-ai collaboration, in: 27th International Conference on Intelligent User Interfaces, IUI '22, ACM, 2022. URL: <http://dx.doi.org/10.1145/3490099.3511135>. doi:10.1145/3490099.3511135.
 - [20] C. Oh, J. Song, J. Choi, S. Kim, S. Lee, B. Suh, I lead, you help but only with enough details: Understanding user experience of co-creation with artificial intelligence, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18, ACM, 2018, p. 1–13. URL: <http://dx.doi.org/10.1145/3173574.3174223>. doi:10.1145/3173574.3174223.
 - [21] C. Zhang, C. Yao, J. Wu, W. Lin, L. Liu, G. Yan, F. Ying, Storydrawer: A child-ai collaborative drawing system to support children's creative visual storytelling, in: CHI Conference on Human Factors in Computing Systems, CHI '22, ACM, 2022, p. 1–15. URL: <http://dx.doi.org/10.1145/3491102.3501914>. doi:10.1145/3491102.3501914.
 - [22] V. N. Antony, C.-M. Huang, Id.8: Co-creating visual stories with generative ai, ACM Transactions on Interactive Intelligent Systems 14 (2024) 1–29. URL: <http://dx.doi.org/10.1145/3672277>. doi:10.1145/3672277.
 - [23] F. Cramer, G. E. Shephard, P. J. Heron, The misuse of colour in science communication, Nature Communications 11 (2020). URL: <http://dx.doi.org/10.1038/s41467-020-19160-7>. doi:10.1038/s41467-020-19160-7.
 - [24] TensorFlow, TensorFlow.js DeepLab Model, <https://github.com/tensorflow/tfjs-models/tree/master/deeplab>, 2024. Accessed: June 5, 2025.
 - [25] OpenCV, OpenCV.js Tutorials, https://docs.opencv.org/4.x/d5/d10/tutorial_js_root.html, 2024. Accessed: June 5, 2025.
 - [26] Y. Ye, F. Sauer, K.-L. Ma, K. Aditya, J. Chen, A user-centered design study in scientific visualization targeting domain experts, IEEE Transactions on Visualization and Computer Graphics 26 (2020) 2192–2203. URL: <http://dx.doi.org/10.1109/TVCG.2020.2970525>. doi:10.1109/tvcg.2020.2970525.