Urban Uchronia: Exploring Alternative Histories and Futures of Cities Using Generative AI and a Projection-Mapped Topography

Marvin Reuter
Technische Hochschule Köln
Gummersbach, Germany
marvin.reuter@th-koeln.de

Tim Yago Nordhoff
Technische Hochschule Köln
Gummersbach, Germany
tim_yago.nordhoff@th-koeln.de

Katrin Schweitzer
Technische Hochschule Köln
Gummersbach, Germany
katrin.schweitzer@th-koeln.de

Sebastian Engelberth
Technische Hochschule Köln
Cologne, Germany
sebastian.engelberth@th-koeln.de

Matthias Böhmer
Technische Hochschule Köln
Cologne, Germany
matthias.boehmer@th-koeln.de

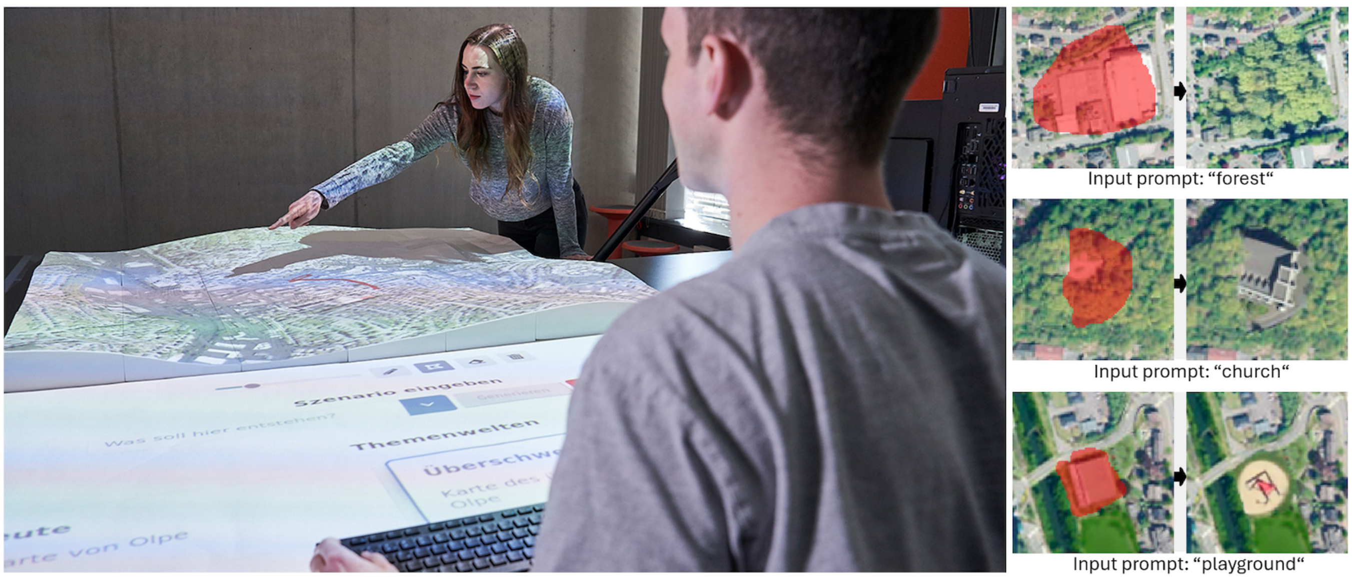

Figure 1: The final prototype in operation, showing the projected satellite imagery on the assembled 3D-printed model (left). Example prompts used to generate alternative places within a city with a forest, a church and a playground by users providing system feedback (right).

Abstract

In museums, visitors are mostly limited to passively observing historical artifacts, with few opportunities for active participation. At the same time, involving citizens meaningfully in urban planning discussions remains a challenge, as the process often feels abstract. Our prototype, *Urban Uchronia*, addresses both challenges by creating an interactive system that combines a physical topographic city model with real-time, AI-generated visualizations projected onto the model, showing alternative urban developments based on the user's imagination. Preliminary feedback from 24 participants indicates that the system has “good” usability (System Usability Scale score: 78.44) and has the potential to stimulate discussions. The prototype establishes a framework that demonstrates the technical feasibility of such systems, its potential to create interactive exhibits, and, as a result, the ability to engage active citizens in urban planning and visualize ideas through textual input.

MuC '25, Chemnitz, Germany

ACM ISBN 979-8-4007-1582-2/25/08
https://doi.org/10.1145/3743049.3748540

CCS Concepts

• **Human-centered computing** → **Interactive systems and tools**; *Geographic visualization*; *Collaborative and social computing devices.*

Keywords

interactive systems, generative AI, topographic city model, projected interface, urban planning, uchronia

ACM Reference Format:
Marvin Reuter, Tim Yago Nordhoff, Katrin Schweitzer, Sebastian Engelberth, and Matthias Böhmer. 2025. Urban Uchronia: Exploring Alternative Histories and Futures of Cities Using Generative AI and a Projection-Mapped Topography. In *Mensch und Computer 2025 (MuC '25), August 31–September 03, 2025, Chemnitz, Germany*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3743049.3748540

1 Introduction and Related Work

Museums increasingly aim to shift from passive forms of presentation towards participatory experiences— an approach equally needed in urban planning, where citizens are often excluded from infrastructural decisions that shape their daily lives, potentially overlooking diverse perspectives and alternative future visions. In both cases, greater public engagement through interactive methods can lead to richer, more forward-thinking outcomes.

Recent research at the intersection of generative AI, urban design, and interactive systems has demonstrated new ways to support public participation and engagement. Several projects have explored how AI-generated imagery can enable non-experts to envision speculative urban futures.

Von Brackel-Schmidt et al. [15] demonstrate how generative AI can support participatory urban planning by enabling citizens to collaboratively visualize speculative future cityscapes using natural language prompts. Preussner et al. [12] extend this idea with UrbAI, a system in which users annotate photographs of urban scenes to receive AI-generated alterations. Both approaches demonstrate the potential of generative systems to lower entry barriers in urban planning contexts. Hung et al. [7] investigate how image-to-image AI can reframe cityscapes in playful or speculative ways, aligning with broader trends in urban storytelling. Further research has demonstrated how generative AI can be used to actively engage residents in urban design by generating realistic images from simple ideas [5, 10].

From a design artifact perspective, Kun et al. [8] present GenFrame, an interactive picture frame that integrates generative AI to produce painterly outputs influenced by user inputs. This work highlights how generative systems can be embedded in traditional cultural artifacts to foster reflection and engagement. In a more environmentally focused artistic context, Canet Sola and Guljajeva [4] explore gaze-based interaction with AI-generated imagery to provoke discourse around environmental destruction, illustrating the narrative power of interactive, generative installations.

Systems such as SandScape [11], HoloCity [9] and CityScope [1] demonstrate the integration of real-time data and urban modeling in AR and tangible interfaces, emphasizing collaborative understanding of complex urban systems. Finally, Günther et al. [6] showcase the use of 3D-printed tangible tools for collaborative interaction on digital surfaces.

Our work extends existing approaches by embedding generative AI within a projection-based interface centered around a physical topographic city model. Unlike prior systems that rely on flat screens or static images, Urban Uchronia enables users to express urban ideas through text and sketches and see their input translate in real time as AI-generated imagery projected directly onto the 3D geography of the city.

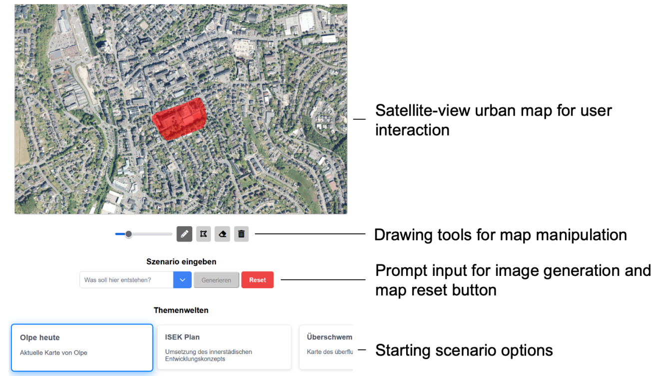

Figure 2: The interface displays an aerial view of the city, with a highlighted area chosen for image generation. Interaction options include sketching, entering text prompts and exploring various alternative or historical maps.

By selecting areas on the projected map and describing potential changes through text prompts, users can generate realistic satellite-style visualizations of alternative urban developments. This creates a hybrid interface that bridges physical spatial reference with digital imagination.

With this paper, we make three main contributions: First, we present the design and implementation of a system that enables intuitive exploration of alternative urban narratives through generative AI. Second, we describe the integration and adaptation of an open-source image generation model specifically fine-tuned for satellite imagery of a German city. Third, we report on the results of initial user feedback, highlighting how citizens engage with our prototype and speculative urban planning in a public setting.

2 Prototype

We developed our prototype for the city of Olpe together with the local museum. Olpe is a mid-sized town located in Germany, NRW, has about 26,000 residents and a history that dates back to AD 900.[1] The concept uses a 3D model of the city's topography, as the geology has played an important role in its historical development, consisting of different hills, valleys, rivers and a dammed-up lake.

2.1 Concept

The core concept behind *Urban Uchronia* is to empower museum visitors to actively shape alternative urban scenarios through intuitive interactions. A hybrid interface combines a physical 3D-printed topographical city model with interactive, AI-generated visualizations projected onto its surface. Users directly mark areas of the projected map and enter textual prompts to generate real-time alternative images of urban development. This interaction aims to encourage visitors to creatively explore speculative urban narratives and reflect on decisions and opportunities in city planning. By allowing users to see the impact of their input projected directly onto familiar geography, the system aims to foster creative speculation and reflective engagement.

[1] https://en.wikipedia.org/wiki/Olpe,_Germany

2.2 Implementation

To ensure smooth operation in a museum context, the software is built as a modular, client-server architecture. It separates visitor interaction from backend services and administrative tools, allowing for flexibility in managing content. This architecture ensures responsive performance, extensibility, and robust operation without usage restrictions. A detailed technical description, including system architecture and technology stack, is provided in the project's open-source repository[2].

2.2.1 Physical Prototype. The physical prototype consists of a 3D-printed topographical model representing the city's urban landscape. Measuring 1.4 × 1.4 meters and assembled from 49 modular tiles, the model provides a spatial reference for user interaction. A vertically mounted projector overlays satellite imagery and AI-generated visualizations directly onto the surface, aligning digital content with the physical geography (see Figure 1).

The model was created using elevation data from the SRTM 90m DEM Digital Elevation Database[3] and scaled to 1:1000 to ensure the visibility of fine-grained urban features such as streets, buildings, and vegetation.

2.2.2 User Interface. Using the mouse and keyboard located next to the model, users can highlight areas of interest through freehand or closed-shape drawing, and enter text prompts. Subsequently, these inputs are used to generate modified satellite-style imagery, which is immediately projected onto the previously selected area of the model (see Figures 1 and 2).

The interface is designed to be intuitive and responsive, supporting both free exploration and guided interaction through curated themes. These include future visions, such as an inner-city development concept; alternative narratives, like a flood simulation; and historical maps from different time periods. It allows users to generate images from different starting scenarios, encouraging imaginative exploration of past, present, or future urban development. An integrated reset function ensures that the system can be returned to a default state to support continuous, unattended use in a public environment. An additional administrative interface allows museum staff to manage maps and configure scenario themes.

2.2.3 AI-Driven Image Generation. At the core of the prototype is an image generation pipeline based on a fine-tuned generative AI model. Once users have selected an area and entered a prompt as inputs, the system generates a realistic visualization. This image is immediately projected onto the corresponding section of the model, maintaining spatial coherence between user intent and visual output. The image generation process is shown in Figure 3.

To support this functionality, we adapted a pre-trained image generation model, Stable Diffusion v2.1 by StabilityAI [13]. The model was fine-tuned using the Diffusers library[4] on a domain-specific dataset of 540 satellite images of a German city and its surrounding areas, which we collected and labeled manually.

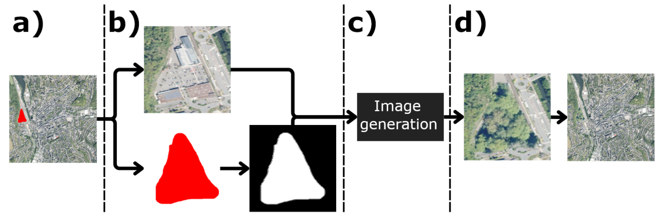

Figure 3: Overview of the image generation process: a) Initial input satellite image with a selected region and the prompt "forest". b) The selected area is cropped, and the image and its corresponding mask are separated. c) A generative AI model uses the prompt, cropped image, and mask to generate the image of a forest. d) The generated image is inserted back into the original satellite image.

The fine-tuned model and training dataset are publicly available on Hugging Face[5]. Three example prompts generated by our fine-tuned AI model are shown in Figure 1, illustrating the model's ability to translate user input into geographically grounded visualizations.

3 User Feedback

To better understand how the prototype is perceived by citizens, we conducted exploratory feedback sessions with two groups of city residents. One group comprises members of the city council with a focus on the fields of education, social work, and sports. The second group consists of pupils of a local high school. Both groups have explored the system for about one hour each.

We aimed to assess the system's usability, to explore how it would be used, whether it fosters engagement with urban planning, and to identify areas for improvement.

To evaluate usability, we employed the German version of the System Usability Scale (SUS) [3, 14], complemented by an open-ended question for additional feedback. We observed how people interacted with the system during the two sessions, noting what emerged. Furthermore, we logged the users' input, including the prompts they entered into the system, the areas they selected on the map, and the resulting images generated.

3.1 Feedback Procedure

The feedback sessions began with a short introduction to the prototype. We explained the system's functionality and demonstrated its main features. Afterward, we motivated the participants to experiment with the prototype. As one person used the prototype directly, others observed and commented, leading to discussions about the prototype's functionality and the proposed urban changes. After usage, the participants completed the feedback questionnaire, which included demographic questions, the SUS-questionnaire, and open-ended feedback.

3.2 Results

In total, 24 participants provided feedback (16 female, 7 male, 1 diverse). Most participants (n=15) were between 16 and 18 years

[2] https://github.com/GP-Alternativweltgeschichten
[3] https://srtm.csi.cgiar.org, accessed on April 27, 2025
[4] https://github.com/huggingface/diffusers
[5] https://huggingface.co/GP-Alternativweltgeschichten

old, while the remaining participants (n=9) ranged from 25 to 76 years (mean=28.25, SD=17.71).

Technical experience was self-assessed on a 5-point Likert scale and distributed evenly across the group. During the study session, 15 participants actively interacted with the prototype, while 9 other participants observed their interactions as bystanders or collaborators without any direct interaction.

The overall System Usability Scale (SUS) score was 78.44 (SD=12.33). Participants who interacted with the prototype directly reported a SUS score of 78.33 (SD=12.81), while observers reported 78.61 (SD=12.25). Despite the age differences, no significant difference in SUS scores was found between the school class (SUS score: 77.81, SD=12.55) and the city council members (SUS score: 79.69, SD=12.64).

In addition to the SUS questionnaire, participants provided open-ended responses. Common themes included requests for higher-quality image generation (i.e., more realistic and detailed visuals), improved usability of interaction tools (i.e., drawing), and suggestions for additional features (i.e., access to more maps and live navigation).

The prompts submitted by participants covered a wide thematic range. Commonly explored topics included natural landscapes (i.e., forests, lakes, rivers), urban infrastructure (i.e., parking areas, parks and playgrounds, residential zones, airports), and historical or cultural landmarks (i.e., castles, churches). Several prompts also focused on natural disasters, imagining scenarios such as hurricanes or volcanic eruptions. In addition, the school class experimented with fictional characters and locations from various contexts, as well as references to well-known public figures, whereas the city council members did not engage with fictional scenarios.

4 Discussion

The system has "good" usability, as indicated by the SUS score of 78.44. Scores above 68 are generally regarded as above average, and scores in the 70–80 range indicate good usability [2]. This suggests that our prototype offers accessible and engaging interactions across different user groups. Consequently, it supports the design goal of enabling unrestricted exploration in an unsupervised museum setting.

The open feedback provides valuable insights into user expectations and highlights areas for improvement. While the core interaction was well received, the feedback regarding image generation fidelity and interface refinement underscores the importance of aligning AI-generated outputs more closely with the users' intentions. Moreover, the interest in expanded functionality points to a strong potential for iterative development beyond the current prototype.

Beyond technical aspects, the collaborations and discussions among participants during sessions indicate that the system does more than visualize changes. The discussions often focused on real issues such as local development, sustainability, and urban identity (e.g., "What if the sports hall was not here, but over there?"). This supports the system's potential not only as a creative tool but also as a medium for dialogues. Hallucination is often a problem in systems using generative AI. However, we observed that in our system, the occurrence of hallucination often provides even more space for speculative discussions. This gives rise to an increased user creativity when interpreting generated visualizations.

At the same time, the system has clear limitations. Despite being custom-trained, the generative model still relies on probabilistic outputs and may not always align with user intentions, especially when prompts are ambiguous. For example, a prompt to generate a festival location for a cultural target-shooting competition (Schützenfest[6]) resulted in a large weapon being placed on the map, which was inappropriate. Such cases highlight the risk of the system generating non-urban objects that may not be child-safe or relevant to the intended purpose of producing satellite-view urban features.

5 Conclusion and Future Work

Urban Uchronia introduced an interactive system that combines a physical topographic model with generative AI to enable collaborative and speculative exploration of alternative urban scenarios. Designed for unsupervised use in public exhibition spaces, the system allows users to mark areas on the model, describe desired changes, and see them reflected in AI-generated satellite-style visualizations.

The system invites active engagement with urban design by making abstract planning ideas understandable, personal, and spatially grounded. Initial feedback indicates *good* usability and potential for reflection and discussion, particularly in settings where public participation and imagination are valued.

We contribute a fully functioning prototype, including the interactive interface, a fine-tuned generative model, and a modular backend, all published as open source for further use and development [7][8].

Future work should further explore and enhance several aspects of our research. First, the quality of AI-generated outputs could be improved by refining the output image scaling, optimizing user input processing, increasing both the size and diversity of the training dataset, and improving the reliability of the AI model. Second, the user interface and the available editing tools could be further refined by enabling selective resetting of AI-generated regions and allowing users to compare the original and modified visualization at any time. Third, the tangibility of the prototype could be enhanced through the integration of 3D-printed tangible tools to support collaborative interaction on the projected visualization, based on the concept demonstrated by Günther et al. [6]. Finally, future research should assess the performance of our system in various domains, including education and urban planning.

Acknowledgments

We thank Sebastian Luke from the Museum of Olpe for his valuable consulting and Jonas Klotzbach, Jan Kolodziejski, Hanna Müller, Collin Seibt, and Ali Ünal for their contributions. This project was funded by the Transfer Fund of TH Köln.

References

[1] Luis Alonso, Yan Ryan Zhang, Arnaud Grignard, Ariel Noyman, Yasushi Sakai, Markus ElKatsha, Ronan Doorley, and Kent Larson. 2018. CityScope: A Data-Driven Interactive Simulation Tool for Urban Design. Use Case Volpe. In *Unifying*

[6] https://en.wikipedia.org/wiki/Sch%C3%BCtzenfest
[7] https://github.com/GP-Alternativweltgeschichten
[8] https://huggingface.co/GP-Alternativweltgeschichten

Themes in Complex Systems IX, Alfredo J. Morales, Carlos Gershenson, Dan Braha, Ali A. Minai, and Yaneer Bar-Yam (Eds.). Springer International Publishing, Cham, 253–261. doi:10.1007/978-3-319-96661-8_27

[2] Aaron Bangor, Philip Kortum, and James Miller. 2009. Determining what individual SUS scores mean: adding an adjective rating scale. *J. Usability Studies* 4, 3 (May 2009), 114–123.

[3] john Brooke. 1996. SUS: A 'Quick and Dirty' Usability Scale. In *Usability Evaluation In Industry*. CRC Press. Num Pages: 6.

[4] Mar Canet Sola and Varvara Guljajeva. 2024. Visions of Destruction: Exploring a Potential of Generative AI in Interactive Art. In *Proceedings of the 17th International Symposium on Visual Information Communication and Interaction (VINCI '24)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3678698.3687185

[5] Jose A. Guridi, Angel Hsing-Chi Hwang, Duarte Santo, Maria Goula, Cristobal Cheyre, Lee Humphreys, and Marco Rangel. 2025. From Fake Perfects to Conversational Imperfects: Exploring Image-Generative AI as a Boundary Object for Participatory Design of Public Spaces. *Proc. ACM Hum.-Comput. Interact.* 9, 2 (May 2025), CSCW014:1–CSCW014:33. doi:10.1145/3710912

[6] Sebastian Günther, Martin Schmitz, Florian Müller, Jan Riemann, and Max Mühlhäuser. 2017. BYO*: Utilizing 3D Printed Tangible Tools for Interaction on Interactive Surfaces. In *Proceedings of the 2017 ACM Workshop on Interacting with Smart Objects*. ACM, Limassol Cyprus, 21–26. doi:10.1145/3038450.3038456

[7] Peng-Kai Hung, Janet Yi-Ching Huang, Stephan Wensveen, and Rung-Huei Liang. 2024. Re.Dis.Cover Place with Generative AI: Exploring the Experience and Design of City Wandering with image-to-image AI. In *Companion Publication of the 2024 ACM Designing Interactive Systems Conference (DIS '24 Companion)*. Association for Computing Machinery, New York, NY, USA, 219–223. doi:10.1145/3656156.3663691

[8] Peter Kun, Matthias Anton Freiberger, Anders Sundnes Løvlie, and Sebastian Risi. 2024. GenFrame – Embedding Generative AI Into Interactive Artifacts. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 714–727. doi:10.1145/3643834.3660750

[9] Oliver Lock, Tomasz Bednarz, and Christopher Pettit. 2019. HoloCity – exploring the use of augmented reality cityscapes for collaborative understanding of high-volume urban sensor data. In *Proceedings of the 17th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry (VRCAI '19)*. Association for Computing Machinery, New York, NY, USA, 1–2. doi:10.1145/3359997.3365734

[10] Yasmin Esmeralda Lützelschwab, Stefan Meier, Florian Mathis, and Adrian Preussner. 2025. CitizenVision: A Mobile App to Promote Citizen Engagement with Generative AI. In *Companion Publication of the 2025 ACM Designing Interactive Systems Conference (DIS '25 Companion)*. Association for Computing Machinery, New York, NY, USA, 554–559. doi:10.1145/3715668.3736341

[11] Ben Piper, Carlo Ratti, and Hiroshi Ishii. 2002. Illuminating clay: a 3-D tangible interface for landscape analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Minneapolis, Minnesota, USA) *(CHI '02)*. Association for Computing Machinery, New York, NY, USA, 355–362. doi:10.1145/503376.503439

[12] Adrian Preussner, Anna Crescenza, Johannes Schöning, and Florian Mathis. 2025. UrbAI: Exploring the Possibilities of Generative AI Image Processing to Promote Citizen Participation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3706598.3713803

[13] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10684–10695. doi:10.48550/arXiv.2112.10752

[14] Bernard Rummel. 2013. System Usability Scale (Translated into German). https://www.researchgate.net/publication/272830038_System_Usability_Scale_Translated_into_German

[15] Constantin von Brackel-Schmidt, Emir Kučević, Stephan Leible, Dejan Simic, Gian-Luca Gücük, and Felix N. Schmidt. 2024. Equipping Participation Formats with Generative AI: A Case Study Predicting the Future of a Metropolitan City in the Year 2040. In *HCI in Business, Government and Organizations*, Fiona Fui-Hoon Nah and Keng Leng Siau (Eds.). Springer Nature Switzerland, Cham, 270–285. doi:10.1007/978-3-031-61315-9_19