Expanding Access to Urban Art at Scale: A Demo of ArtScape

Lucy Jiang*
Human Centered Design and Engineering
University of Washington
Seattle, Washington, USA
lucjia@uw.edu

Jared Hwang*
Paul G. Allen School of Computer Science & Engineering
University of Washington
Seattle, Washington, USA
jaredhwa@cs.washington.edu

Jon E. Froehlich
Paul G. Allen School of Computer Science & Engineering
University of Washington
Seattle, Washington, USA
jonf@cs.uw.edu

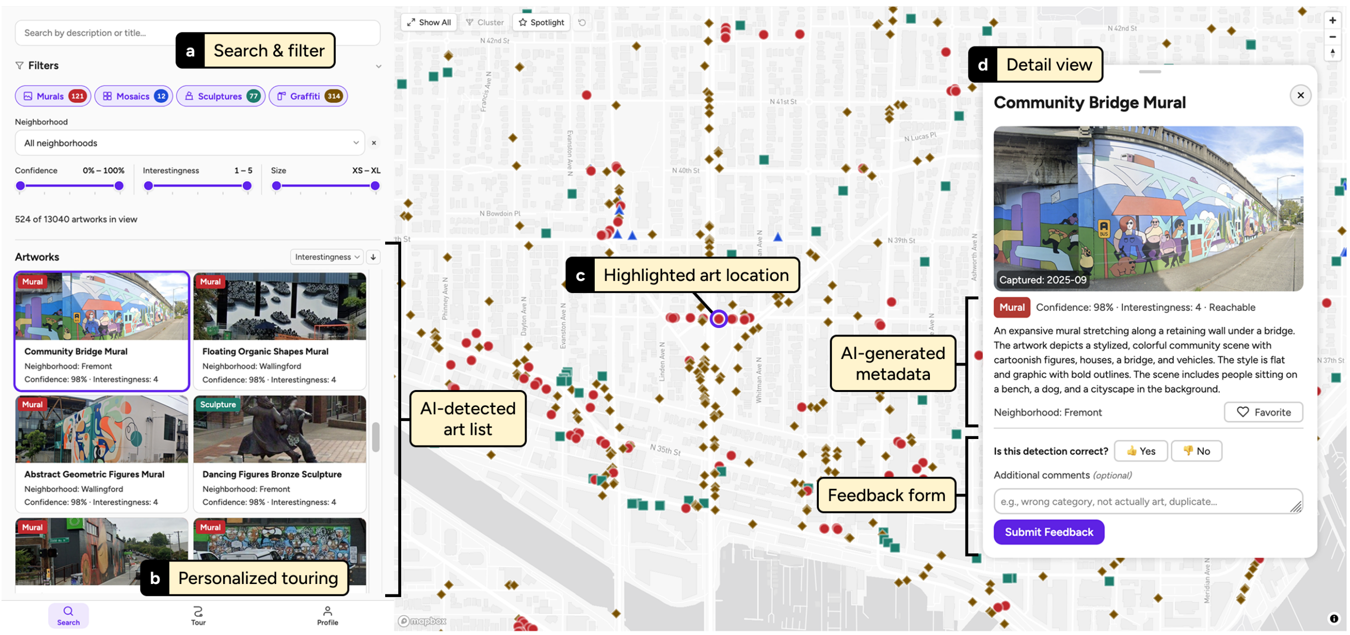

Figure 1: We present ArtScape, an accessible web tool that allows users to search for artwork, filter by type or attribute, and generate tours to explore artwork along a specific route, all through a screen reader accessible interface.

Abstract

Public art fundamentally shapes urban environments, fostering community identity and providing cultural and aesthetic enrichment. However, most urban artwork is inaccessible to blind and low vision (BLV) people. We introduce *ArtScape*, an urban art exploration prototype that allows users to discover, inspect, and explore urban artwork throughout a city, powered by a novel VLM-based pipeline that automatically detects and characterizes artwork from street-view imagery. We present the design of ArtScape and findings from a user study evaluation including three BLV participants.

CCS Concepts

• **Human-centered computing** → **Accessibility**; **Empirical studies in accessibility**.

*Both authors contributed equally to this research.

ACM ISBN 979-8-4007-2521-0/2026/10
https://doi.org/10.1145/3797867.3841240

Keywords

urban art, blind, low vision, urban accessibility, art accessibility, AI

ACM Reference Format:
Lucy Jiang, Jared Hwang, and Jon E. Froehlich. 2026. Expanding Access to Urban Art at Scale: A Demo of ArtScape. In *The 28th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '26), October 25–28, 2026, Vila Nova de Gaia, Portugal.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3797867.3841240

1 Introduction

From murals to street installations to graffiti, urban art embodies the cultural, social, and political identity of neighborhoods. Yet, most people have no systematic way to discover or learn about artwork in their cities unless local governments or communities maintain dedicated indices. This gap is particularly acute for blind and low vision (BLV) people, as there are limited accessible ways to engage with and understand the public art that surrounds them in their daily lives or when touring new cities. Existing research on detecting artwork has focused on detecting single types of art, such as graffiti [2, 6, 9, 11, 19, 21, 29, 30, 32], murals [2, 11, 33] and sculptures [13, 24], but makes no attempt to *characterize* or *describe*

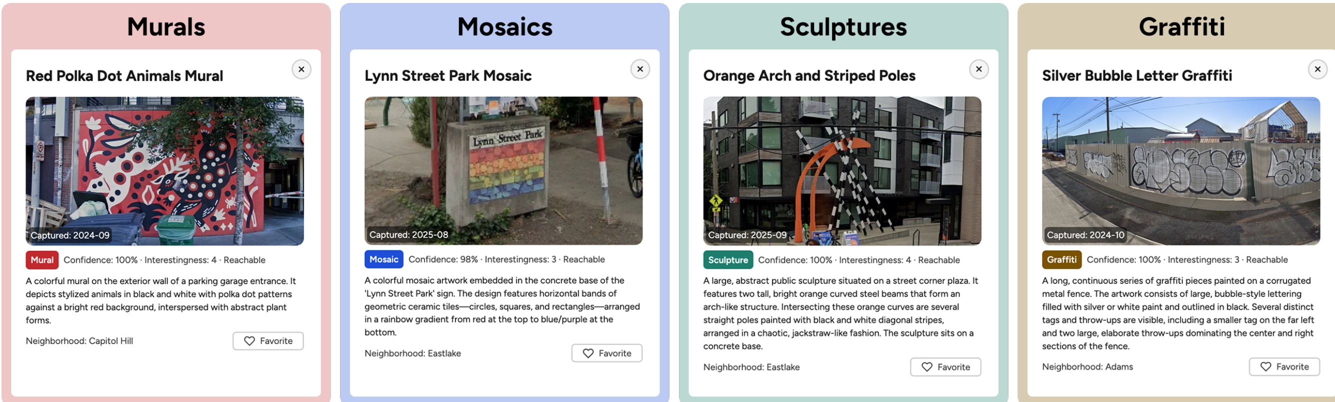

Figure 2: ArtScape identifies art in four categories: murals, mosaics, sculptures, and graffiti.

the artwork accessibly, despite other literature establishing the efficacy for AI to generate accessible artwork descriptions [3, 10, 17]. Furthermore, existing curated artwork lists are largely manually created [23, 25, 26, 31], and self-driven, personalized artwork discovery is generally not supported [8, 22, 28].

However, through advances in streetscape analysis and vision language models (VLMs)—designed to navigate subjectivity via human language processing—we can scalably identify and describe urban artwork from widely available ground-level imagery. In this demo paper, we introduce ArtScape, the first automated, large-scale urban artwork exploration tool that detects and describes urban artwork (*murals*, *mosaics*, *sculptures*, and *graffiti*) in Google Street View panoramas and interactively presents these detections in a screen-reader accessible interface. To evaluate ArtScape, we conducted a 12-participant user study, including three blind or low vision participants. We found that this digital urban art catalog improved accessibility for both BLV and sighted participants and was a positive application of AI for artwork and cultural understanding, though BLV participants were generally less confident about ArtScape's artwork detection comprehensiveness compared to sighted participants. Key challenges included noise introduced by false positives, which were somewhat difficult for BLV people to identify, and limited cultural context in the visual descriptions.

During the ASSETS 2026 demo session, we will showcase the ArtScape system and allow users to interact with and explore the urban artwork our system has detected in Seattle, WA. Attendees can interact with ArtScape either visually or with a screen reader.

Overall, our work advances research in scalable streetscape characterization and increases accessibility to a previously inaccessible, spontaneous feature of the built environment. We contribute insights on how VLM-powered methods can support BLV people's access to urban artwork and identify areas for future work. Though the design and development of ArtScape was motivated and informed by needs identified in prior work on urban art accessibility [15, 17], we believe its broad utility for both sighted and BLV people demonstrates a positive "curb cut effect" of making urban art more accessible.

2 The ArtScape System

We present ArtScape, the first screen-reader accessible urban artwork cataloging and exploration tool.

2.1 Defining Urban Artwork

We ground ArtScape's typology of artwork in existing urban artwork literature. While prior work has stated that urban artwork must be socially engaged [18], non-commercial [4], site-specific [18], and environmentally integrated [12], we define urban artwork as any artwork that is outdoors and publicly accessible. This builds on Jiang *et al.*'s urban artwork taxonomy which consists of murals (2D, on flat surfaces), mosaics (2.5D, also on flat surfaces), and sculptures (3D, typically freestanding or on pedestals) [15, 17]. We additionally include graffiti in our typology due to its ubiquity and cultural meaning. Figure 2 shows examples of detections for each of our four artwork categories.

2.2 Detecting Urban Artwork

ArtScape is powered by a five-stage VLM-powered pipeline to find, categorize, and describe urban artwork. (1) We begin by sampling and downloading Google Street View panoramas at approximately every 30 meters. (2) Then, we prompt Gemini 3 Pro [1] to find every instance of urban artwork in the panorama, categorize the artwork, and generate descriptions of each artwork. (3) We pass crops of the artwork detected in stage 2 to a second pass through the VLM to both remove false positives and generate more accurate descriptions. (4) We then use a depth estimation model, DA360 [14], to geolocate each art piece in the real world. (5) Lastly, we deduplicate artwork detections of the same piece by generating embedding vectors using DINOv3 [27] of each crop and removing pieces with high cosine similarity (above 0.88).

2.3 Exploring Urban Artwork

ArtScape includes four key components for viewing and inspecting urban artwork: (1) a search panel that displays artwork as cards in a structured list that allows users to filter by text or by category; (2) a tour functionality that generates routes that pass by artworks of interest; (3) a map view that displays all artworks at their detected location; and (4) a detail view that provides more information

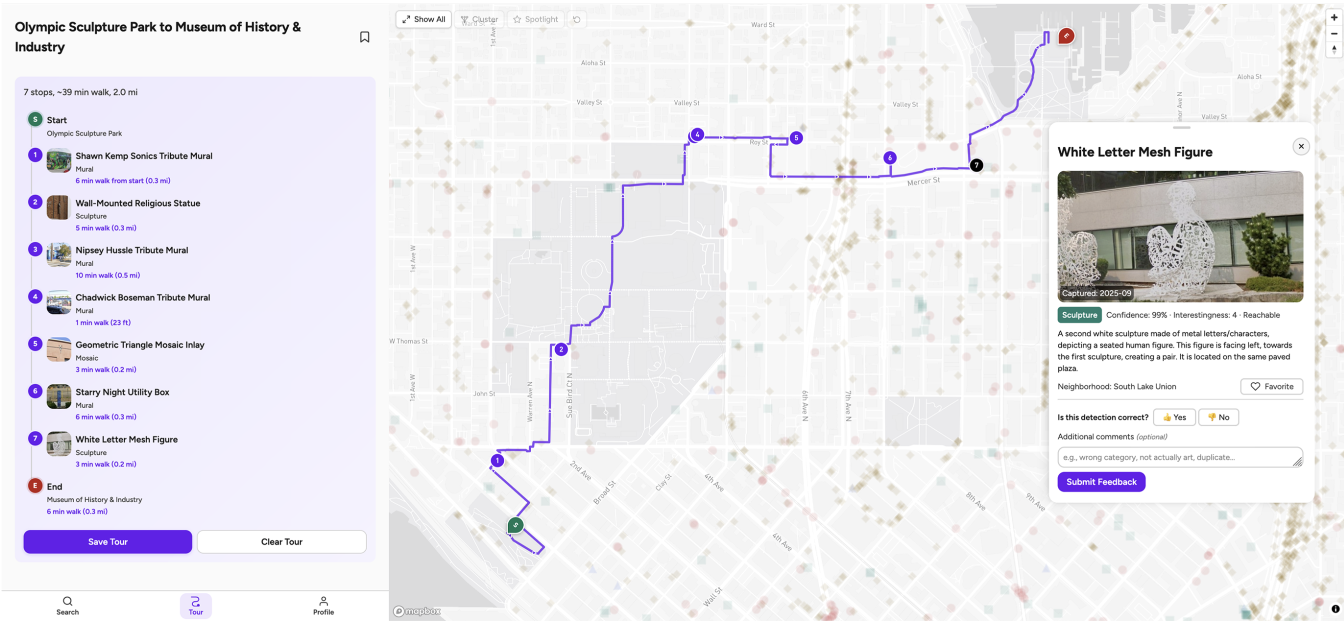

Figure 3: ArtScape includes a tour functionality that allows users to specify an origin, destination, number of stops, and types of artwork, then generates a customized artwork tour.

about a particular art piece. Figure 1 labels key components of the ArtScape system and Figure 3 shows an example artwork tour.

We designed our web application to be screen reader accessible and usable by low vision users by providing nonvisual information for visual elements on the interface and considering the cognitive load required for engaging with the database. For example, though sighted people can visually perceive artwork density from clustered map markers, we list the number of artworks in each neighborhood as part of the dropdown menu to make the density more accessible to BLV users. In addition to visually displaying artwork descriptions on individual cards, which makes these descriptions legible to sighted users, we also generate and embed separate alt text that is shorter and more concise to avoid the "screen reader stutter" effect [7]. To support clarity and reduce redundancy with two different description lengths, on the detail view, we prepend "short description" to the alt text and indicate that the long description is visible on the page to highlight that screen reader users can receive a more detailed description if they wish (Figure 4). We also use different icons and colors to indicate each artwork type on the map, making it more accessible to people with color vision deficiency or low vision rather than relying on color alone.

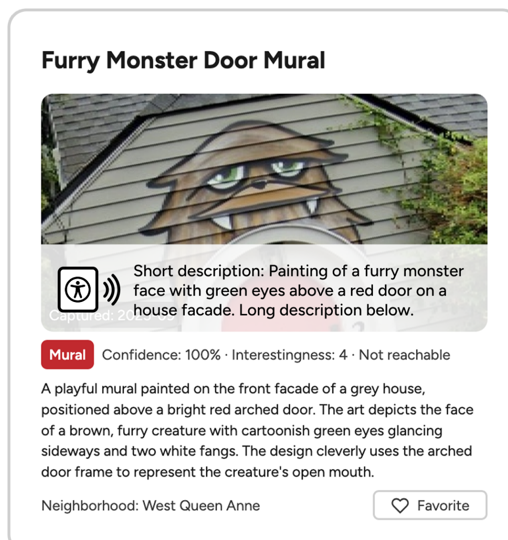

Figure 4: Each detail view pairs a short alt text (only exposed to screen readers; overlaid on the image for illustration) with a longer description that is visible on the page.

3 User Study

3.1 Method

To evaluate ArtScape, we conducted virtual interviews with 12 participants, including three BLV people (B#), three sighted artists (A#), two sighted urban planners (U#), and four sighted people (S#). We interviewed participants with varied expertise to explore how the system could improve accessibility both in terms of nonvisual access, motivated by prior work [17], but also logistical and intellectual access. During our study, we first asked participants about their urban artwork experiences, then engaged participants in a five-task prototype evaluation where they explored artwork at the aggregate and individual level, planned an urban art walking tour, and assessed AI-generated information. We closed with a reflection on their prototype experience and questions about what participants considered to be urban artwork. Two authors analyzed the interview recordings using a hybrid inductive-deductive coding method, with attention to both similarities and differences across the diverse stakeholder groups. In this demo paper, we focus on accessibility-oriented findings from our BLV participants, centering their deep expertise with visual descriptions and their perspectives on AI-mediated access. Findings from non-BLV participants can be found in our full paper [16].

3.2 Findings

ArtScape dramatically improved the accessibility of exploring urban artwork, with all BLV participants raising their ratings from 1 to 5. While some BLV participants highlighted that art exploration was much more accessible with a sighted companion, they felt that ArtScape enabled more independence: *"using this tool, I would give it a 5 [...] without this tool, it's basically non-existent"* (B6). Many of the other participants highlighted the convenience and usefulness of having a searchable repository of geolocated urban artwork, also describing the tool as *"making urban artwork a lot more accessible to find"* (U7). B4 specifically highlighted the curb cut effect of ArtScape and how it could *"open up artwork to all kinds of people, not just disabled people [...] I would say most people don't know the plethora of artwork in their neighborhood [or] city."* Beyond locating and describing artwork, BLV participants also valued that ArtScape distinguished different artwork types, given that some types of artwork are more accessible than others (*e.g.*, sculptures are more tactile than murals). In particular, B1 shared that the tool could be especially helpful for finding sculptures *"if I really want to actually get out there and touch things."*

Furthermore, all participants were positive about the usage of AI for detecting and describing urban artwork at scale, but they also noted limitations. In contrast to *"controversy around AI-generated art"* (S3), participants felt as though using VLMs for detecting artwork that already existed was an appropriate and favorable application. Even then, participants acknowledged some drawbacks of this AI-based approach, such as featuring many false positive detections and limited cultural context in the visual descriptions.

Though all participants appreciated ArtScape and the insights it unlocked, BLV and sighted participants varied in their ratings of how comprehensive they believed the ArtScape detections were. Sighted participants rated comprehensiveness as 4.28/5, citing the sheer volume and density of detections as well as false positives they noticed (*e.g.*, small graffiti tags on bicycles) as justification. On the other hand, BLV participants gave an average rating of 3.17/5, which could be attributed to their limited insight into the density of detections, difficulties of nonvisually identifying false positives, and their prior poor experiences with AI visual identification systems leading them to have lower expectations: *"I don't have any reason to feel confident"* (B6).

Beyond the existing features—including the AI-generated short and long descriptions—BLV participants gave additional suggestions to improve the overall accessibility of ArtScape. For example, providing the precise address of the urban artwork, as well as turn-by-turn directions, could improve the accessibility of the tour functionality. B1 shared his curiosity about whether a detection was *"really a mural, or just a picture in front of a restaurant,"* suggesting that additional context could also support their evaluation of whether a detection was a false positive.

4 Discussion

In this demo, we introduced ArtScape, a large-scale urban artwork exploration tool that uses VLMs to detect and describe urban artwork from Google Street View images. Below, we discuss the usage of AI for accessibility and opportunities for future HCI work to further improve the accessibility of urban artwork.

4.1 Urban Art Access

Both BLV and sighted participants felt that ArtScape improved their access to urban artwork, demonstrating how designing and developing systems for nonvisual access can improve outcomes for all (*i.e.*, the curb cut effect [5]). Furthermore, given the distributed and decentralized nature of urban artwork, ArtScape not only provided visual descriptions to support access for BLV people but also improved accessibility in other ways by making it possible to engage with public art from any location and at any time. By making urban art more accessible, ArtScape further enhances access to the cultural, social, and political conversations associated with it.

4.2 Limitations and Future Work

We propose the following areas of future work. First, many sighted people expressed that additional information would be helpful in the descriptions—future research could enrich the artwork descriptions to include more cultural context of each art piece (*e.g.*, the artist, historical background). Additionally, to support in-situ urban artwork experiences, a mobile application could simulate a sighted user's experience of spontaneous artwork discovery for a BLV user by notifying them of artwork in their immediate vicinity. To improve nonvisual access to the map featuring all of the artwork detections, research could incorporate emerging work (*e.g.*, [20]) on accessible data visualizations to help BLV users observe regional trends in artwork more easily.

We also recognize that our evaluation with blind and low vision users was somewhat limited. ArtScape addresses the challenge of urban artwork scale, one of multiple needs identified through prior work focusing on BLV people's preferences for urban art access [17], but we acknowledge that descriptions are but one way to interact with artwork. We recommend future work to explore how to make VLM-identified artwork accessible in a variety of modalities, including but not limited to tactile and auditory methods.

5 Conclusion

Our work introduces the first accessible and scalable urban artwork detection and exploration platform, increasing access to a highly visual yet inaccessible form of artwork and expression in our built environments. We also contribute insights on BLV users' perspectives on AI-mediated urban art access at scale. Our ASSETS 2026 demo invites attendees to use ArtScape, either visually or with a screen reader, to explore Seattle's urban artwork.

Acknowledgments

We thank all study participants for their time. This work was supported by NSF grant #2411222 and partially supported by Google Cloud Credits from Google.org.

References

[1] 2025. Gemini 3.1 Pro. https://deepmind.google/models/gemini/pro/.

[2] Javier Rozo Alzate, Marta S. Tabares, and Paola Vallejo. 2021. Graffiti and Government in Smart Cities: A Deep Learning Approach Applied to Medellín City, Colombia. In *International Conference on Data Science, E-learning and Information Systems 2021 (DATA'21)*. Association for Computing Machinery, New York, NY, USA, 160–165. doi:10.1145/3460620.3460749

[3] Takaya Arita, Wenxian Zheng, Reiji Suzuki, and Fuminori Akiba. 2025. Assessing LLMs in Art Contexts: Critique Generation and Theory of Mind Evaluation. arXiv:2504.12805 [cs] doi:10.48550/arXiv.2504.12805

[4] Andrea Lorenzo Baldini. 2022. What is street art? *Estetika: The European Journal of Aesthetics* 59, 1 (2022), 1–21. https://doi.org/10.33134/eeja.234

[5] Angela Glover Blackwell. 2017. *The Curb-Cut Effect.* https://ssir.org/articles/entry/the_curb_cut_effect

[6] Tácio Souza Bomfim, Éldman de Oliveira Nunes, A. Belén Moreno, and Ángel Sánchez. 2025. Graffiti and Covered Material Detections in Urban Images Using Deep Learning. In *2025 32nd International Conference on Systems, Signals and Image Processing (IWSSIP)*. 1–6. doi:10.1109/IWSSIP66997.2025.11151881

[7] Sheri Byrne-Haber. 2026. The Screen Reader Stutter: Why Your Content Repeats Itself and How to Fix It. https://www.sheribyrnehaber.com/the-screen-reader-stutter-why-your-content-repeats-itself-and-how-to-fix-it/

[8] Ricardo Campos and Ágata Sequeira. 2020. Urban art touristification: the case of Lisbon. *Tourist Studies* 20, 2 (2020), 182–202. https://doi.org/10.1177/1468797619873108

[9] Patrikakis Charalampos, Kasnesis Panagiotis, Toumanidis Lazaros, and Tzitamidis Alexandros. 2019. STORM Graffiti/Tagging Detection Dataset. doi:10.5281/zenodo.3238357

[10] Arnavi Chheda-Kothary, Ritesh Kanchi, Chris Sanders, Kevin Xiao, Aditya Sengupta, Melanie Kneitmix, Jacob O. Wobbrock, and Jon E. Froehlich. 2025. ArtInsight: Enabling AI-Powered Artwork Engagement for Mixed Visual-Ability Families. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25)*. Association for Computing Machinery, New York, NY, USA, 190–210. doi:10.1145/3708359.3712082

[11] Joana Fogaça, Tomás Brandão, and João C. Ferreira. 2023. Deep Learning-Based Graffiti Detection: A Study Using Images from the Streets of Lisbon. *Applied Sciences* 13, 4 (Jan. 2023), 2249. doi:10.3390/app13042249

[12] Hilde Hein. 2006. *Public art: Thinking museums differently.* Bloomsbury Publishing PLC.

[13] Dajeong Hong and Jongweon Kim. 2017. Sculpture Detection Method Using the Convolution Neural Network. In *Proceedings of the 2017 International Conference on Information Technology (ICIT '17)*. Association for Computing Machinery, New York, NY, USA, 148–151. doi:10.1145/3176653.3176727

[14] Hualie Jiang, Ziyang Song, Zhiqiang Lou, Rui Xu, and Minglang Tan. 2025. Depth Anything in 360°: Towards Scale Invariance in the Wild. arXiv:2512.22819 [cs] doi:10.48550/arXiv.2512.22819

[15] Lucy Jiang, Jon E Froehlich, and Leah Findlater. 2024. Making Urban Art Accessible: Current Art Access Techniques, Design Considerations, and the Role of AI. *arXiv preprint arXiv:2410.20571* (2024). https://doi.org/10.48550/arXiv.2410.20571

[16] Lucy Jiang, Jared Hwang, and Jon E. Froehlich. 2026. ArtScape: Detecting, Characterizing, and Interactively Exploring Urban Art at Scale Using Vision Language Models. In *Proceedings of the 39th Annual ACM Symposium on User Interface Software and Technology*. 1–19. https://doi.org/10.1145/3830398.3830557

[17] Lucy Jiang, Amy Seunghyun Lee, Jon E. Froehlich, and Leah Findlater. 2026. Unseen City Canvases: Exploring Blind and Low Vision People's Perspectives on Urban and Public Art Accessibility. https://doi.org/10.48550/arXiv.2603.26909

[18] Suzanne Lacy. 1995. Mapping the terrain: New genre public art. (1995). https://monoskop.org/images/7/7c/Lacy_Suzanne_ed_Mapping_the_Terrain_New_Genre_Public_Art_1995.pdf

[19] Bahram Lavi, Eric K. Tokuda, Felipe Moreno-Vera, Luis Gustavo Nonato, Claudio T. Silva, and Jorge Poco. 2022. 17K-Graffiti: Spatial and Crime Data Assessments in São Paulo City. In *17th International Conference on Computer Vision Theory and Applications*. 968–975.

[20] Chu Li, Rock Yuren Pang, Arnavi Chheda-Kothary, Ather Sharif, Henok Assalif, Jeffrey Heer, and Jon E. Froehlich. 2026. GeoVisA11y: An AI-based Geovisualization Question-Answering System for Screen-Reader Users. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–17. https://doi.org/10.1145/3772318.3790334

[21] Tessio Novack, Leonard Vorbeck, Heinrich Lorei, and Alexander Zipf. 2020. Towards Detecting Building Facades with Graffiti Artwork Based on Street View Images. *ISPRS International Journal of Geo-Information* 9, 2 (Feb. 2020), 98. doi:10.3390/ijgi9020098

[22] Punartha Jeevana Perera. 2019. Urban art scene in Madrid: How can contemporary art be used for tourism? *Enlightening Tourism. A Pathmaking Journal* 9, 1 (2019), 1–37. https://doi.org/10.33776/et.v9i1.3552

[23] Preservation Detroit. 2026. *Public Art in the Cultural Center Tour.* https://preservationdetroit.org/public-art-in-the-cultural-center-tour

[24] Mubarak Auwalu Saleh, Zubaida Said Ameen, Chadi Altrjman, and Fadi Al-Turjman. 2022. Computer-Vision-Based Statue Detection with Gaussian Smoothing Filter and EfficientDet. *Sustainability* 14, 18 (Jan. 2022), 11413. doi:10.3390/su141811413

[25] Seattle Center. 2026. *Public Art.* https://www.seattlecenter.com/explore/arts/public-art

[26] Seattle Office of Arts & Culture. 2026. *Works.* https://seattlearts.emuseum.com/objects

[27] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. 2025. DINOv3. arXiv:2508.10104 [cs] doi:10.48550/arXiv.2508.10104

[28] Jonathan Skinner and Lee Jolliffe. 2017. 'Wall-to-wall coverage': An introduction to murals tourism. *Murals and tourism* (2017), 1–24. https://doi.org/10.4324/9781315547978-1

[29] Eric K. Tokuda, Roberto M. Cesar, and Claudio T. Silva. 2019. Quantifying the Presence of Graffiti in Urban Environments. In *2019 IEEE International Conference on Big Data and Smart Computing (BigComp)*. 1–4. doi:10.1109/BIGCOMP.2019.8679113

[30] Andrea Vallebueno and Yong Suk Lee. 2023. Measuring Urban Quality and Change through the Detection of Physical Attributes of Decay. *Scientific Reports* 13, 1 (Oct. 2023), 17316. doi:10.1038/s41598-023-44551-3

[31] Visit Detroit. 2026. *Detroit Mural Guide.* https://visitdetroit.com/inside-the-d/detroit-mural-guide/

[32] Bingshu Wang, Qianchen Mao, Aifei Liu, Long Chen, and C. L. Philip Chen. 2025. DIG: Improved DINO for Graffiti Detection. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 55, 5 (May 2025), 3557–3569. doi:10.1109/TSMC.2025.3541795

[33] Penglei Wang, Xin Fan, Qimeng Yang, Shengwei Tian, and Long Yu. 2025. Object Detection of Mural Images Based on Improved YOLOv8. *Multimedia Systems* 31, 2 (Feb. 2025), 93. doi:10.1007/s00530-025-01687-8