
Sound Clouds: Exploring ambient intelligence in public spaces to elicit deep human experience of awe, wonder, and beauty

Chengzhi Zhang¹ Dashiell Carrera^{1,2,†} Daksh Kapoor^{1,†} Jasmine Kaur^{1,†} Jisu Kim^{1,†}
Brian Magerko¹

¹Georgia Institute of Technology, Atlanta, GA 30309, USA

²University of Toronto, Toronto, ON, Canada

¹{czhang694, dkapoor30, jkaur47, jisu.kim, magerko}@gatech.edu,

²dcarrera@dgp.toronto.edu

[†]Equal contribution



Figure 1: Participants Interacting with *Sound Clouds* on the Live Exhibition Day

Abstract

While the ambient intelligence (AmI) systems we encounter in our daily lives, including security monitoring and energy-saving systems, typically serve pragmatic purposes, we wonder how we can design and implement ambient artificial intelligence experiences in public spaces that elicit deep human feelings of awe, wonder, and beauty. As a manifestation, we introduce—*Sound Clouds*, an immersive art installation that generates live music based on participants’ interaction with several human-height spheres. Our installation serves as a provocation into future ambient intelligence that provokes, not limits, the future possibilities of AmI.

1 Introduction

Ambient intelligence (AmI), as a subset of artificial intelligence, has started to play an increasingly inseparable role in our daily technological encounters. While there are varied definitions of AmI, one widely adopted definition from Aarts for ambient intelligence is “developing technology that will increasingly make our everyday environment sensitive and responsive to our presence” Aarts and Encarnação [2006]. AmI generally possesses the characteristics of being Sensitive (S), Responsive (R), Adaptive (A), Transparent (T), Ubiquitous (U), and Intelligent (I) Cook et al. [2009]. While most of the ambient intelligence serves a pragmatic purpose, such as a) the *Nest thermostat* Hernandez et al. [2014] for achieving domestic energy-saving, b) *ambient influence* for promoting positive behavior change like taking the stairs Rogers et al. [2010], and c) *LAICA* for enhancing public security Cucchiara et al. [2005], etc., there has been little exploration in other kinds of human experience with ambient intelligence. Therefore, we ask the following overarching research question: *How can we design interactions with ambient artificial intelligence in public spaces to elicit deep human feelings of awe, wonder, and beauty?*

We introduce *Sound Clouds*, an exploratory ambient intelligence art installation that consists of several larger-than-human-height (8-12' in diameter) floating spheres that were filled with helium and air to (approximately) reach neutral buoyancy. Participants can use any or all of their bodies to interact with the spheres while they control real-time music that is generated from the movement and location of the spheres. There is no instruction or explanation of the system given to participants—curiosity alone guides their interactions and exploration.

2 *Sound Clouds*

Sound Clouds is an interactive ambient intelligence environment comprising several human-height Polyvinyl Chloride (PVC) spheres (with diameters from 1.2 m to 3 m). It was exhibited inside the building of a renovated factory with interior dimensions of 18 m (width) \times 42 m (length) \times 9 m (height). *Sound Clouds* was experienced by more than 200 participants during our two-hour public exhibition. Inside the building, we used dry ice and fabric canopies hanging on top of the ceiling to create a serene feeling and projected “waves” of light on the ground to create an underwater theme. We mounted a GoPro¹ camera on the ceiling facing top-down to capture the live video input of the space. The video feed is routed via Wi-Fi to a central-processing Mac computer, which runs a computer vision (CV) model for object detection to obtain spheres’ location and diameter. The diameter is used to determine the sphere’s height. The central-processing Mac computer used the ESP-NOW protocol² to send the light signal to each ESP32 board embedded inside the sphere. According to the spheres’ locations, *Sound Clouds* generates live music in Max/MSP (see Figure 2) software and changes the light of the spheres accordingly. We allowed at most 20 participants to interact simultaneously.




Figure 2: Screenshot of the *Sound Clouds* Music Generation Logic Implemented in Max/MSP

2.1 The roles of AI and ML

We used computer vision (CV) for tracking spheres using a fine-tuned YOLO12n model³. The music generation is rule-based through Max/MSP⁴ software. Although the CV model tracked all the spheres, we used the parameters of two spheres for music generation. We looped an ocean-like sound as the background, depicting the ethereal, mysterious atmosphere. We divided the video feed space into nine zones and used three key scripts for sound generation (see the description below and Figure 3, Right). Due to the fish’s-eye view of the GoPro wide-angle lens, the vision is lightly distorted (see Figure 3, Left). The key scripts for interaction are:

- **cloudGrid**: One sphere enters a new zone, triggering a corresponding and distinguished pre-recorded sound.
- **orbProxi**: Two spheres close enough to each other below a threshold trigger a sound.
- **pitchShift**: The depth of one sphere shifts the pitch of the background sound.

3 *Sound Clouds* and Humanity

While the ambient intelligence usually serves a pragmatic purpose, *Sound Clouds* elicits deep human experience of awe, wonder, and beauty. Awe, as a specific self-transcendent emotion, is characterized

¹GoPro <https://gopro.com/en/us>

²ESP-NOW <https://github.com/espressif/esp-now>

³YOLO12n <https://docs.ultralytics.com/models/yolo12/>

⁴Max/MSP Wikipedia [https://en.wikipedia.org/wiki/Max_\(software\)](https://en.wikipedia.org/wiki/Max_(software))

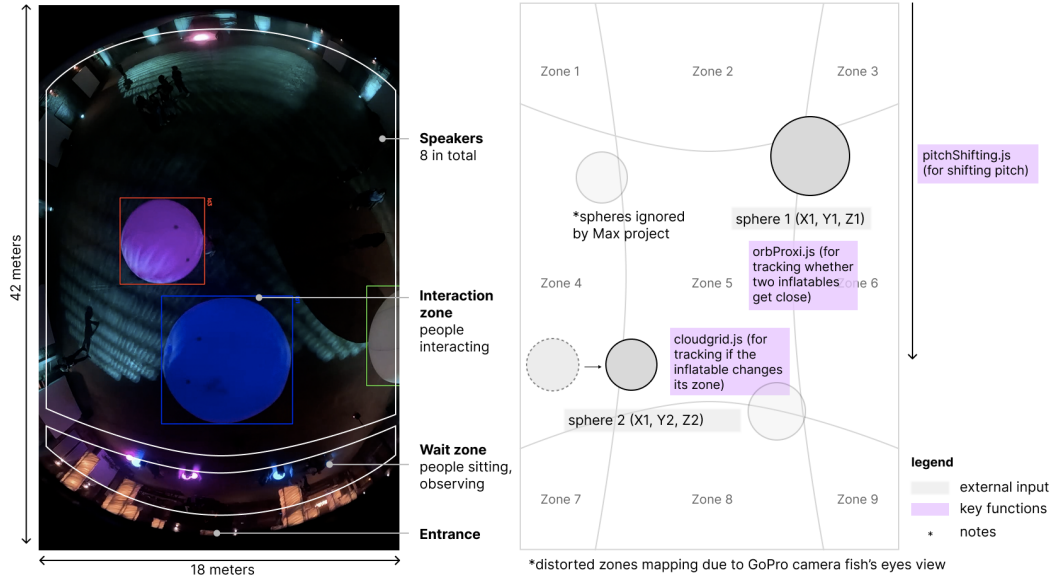


Figure 3: **Left:** Top-down View of the Exhibit Space (View from the Go Pro Camera), **Right:** Top-down Schematic of the Max Music Generation Project

by “the perception of being in the presence of something vast that the individual does not immediately understand” Stellar et al. [2017]. As a rarely-experienced emotion, experiencing awe has been found to prompt feelings of interconnectedness Shiota et al. [2007], motivate commitment to social collectives Keltner and Haidt [2003], and improve perception of well-being Rudd et al. [2012]. Thus, experiencing awe can contribute to the exploration of humanity aspects of AI. Participants shared their feelings of awe during the post-interaction interview, including:

“...the ambience and the whole vibe of the place was very **peaceful and calming**”
 “just like, jaw on the floor...it felt almost like **dream-like**”
 “I think that that’s one of the things that we lose in childhood, from childhood, as adults, we don’t play because the proportions are off. That **sense of wonder and awe** are different.”

Following the research thread of *The Aware Home* Abowd et al. [2002], an AmI research space supporting the elderly’s dignified aging, and *Smart Kitchen* Bell and Kaye [2002]—which prioritizes “experience, affect, and desire” but not just efficiency, our exhibit serves as a provocation for future AmI systems that enrich—not replace—humanity.

4 Author Bibliography

- Chengzhi Zhang is a PhD student at Georgia Tech with a research focus on AI literacy Zhang and Magerko [2025] and the intersection of AI and UX Zhang et al. [2023].
- Dashiel Carrera is a Visiting PhD Researcher at Columbia University, a novelist, and a sound media artist whose research and art practice explore AI’s impact on the Arts Carrera et al. [2025].
- Daksh Kapoor is a graduate student at Georgia Tech with a research focus on embodied interaction and experience design, with an emphasis on how users interpret, engage with, and move within immersive and interactive systems.
- Jasmine Kaur is a graduate student at Georgia Tech with a research focus on human-computer interaction and speculative design with a special emphasis on creative AI technologies Kaur et al. [2025].
- Jisu Kim is an M.S. student and Graduate Research Assistant at Georgia Tech with a research focus on responsible AI Lee et al. [2024] and human-AI collaboration, specifically how AI can augment—rather than replace—human values.

- Brian Magerko is a Regent’s Professor of Digital Media and Director of Graduate Studies at Georgia Tech and has contributed broadly over 3 decades to creativity, cognition, and computing research with a focus on human/AI co-creativity.

References

- E. H. Aarts and J. L. Encarnação. *True visions: The emergence of ambient intelligence*. Springer Science & Business Media, 2006.
- G. D. Abowd, A. F. Bobick, I. A. Essa, E. D. Mynatt, and W. A. Rogers. The aware home: A living laboratory for technologies for successful aging. In *Proceedings of the AAAI-02 workshop “automation as caregiver*, pages 1–7, 2002.
- G. Bell and J. Kaye. Designing technology for domestic spaces: A kitchen manifesto. *Gastronomica*, 2(2):46–62, 2002.
- D. Carrera, Z. Zhao, A. Ajin Thomas, and D. Wigdor. Nabokov’s cards: An ai assisted prewriting system to support bottom-up creative writing. In *Proceedings of the 2025 Conference on Creativity and Cognition*, C&C ’25, page 546–559, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400712890. doi: 10.1145/3698061.3726916. URL <https://doi.org/10.1145/3698061.3726916>.
- D. J. Cook, J. C. Augusto, and V. R. Jakkula. Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and mobile computing*, 5(4):277–298, 2009.
- R. Cucchiara, A. Prati, and R. Vezzani. Ambient intelligence for security in public parks: the laica project. In *IEE International Symposium on Imaging for Crime Detection and Prevention (ICDP 2005)*, pages 139–143. IET, 2005.
- G. Hernandez, O. Arias, D. Buentello, and Y. Jin. Smart nest thermostat: A smart spy in your home. *Black Hat USA*, 2015, 2014.
- J. Kaur, M. Trajkova, A. Knowlton, and B. Magerko. Bringing LuminAI to Life: Studying Dancers’ Perceptions of a Co-Creative AI in Dance Improvisation Class. In *Proceedings of the 2025 Conference on Creativity and Cognition*, pages 859–871, Virtual United Kingdom, June 2025. ACM. doi: 10.1145/3698061.3726925. URL <https://dl.acm.org/doi/10.1145/3698061.3726925>.
- D. Keltner and J. Haidt. Approaching awe, a moral, spiritual, and aesthetic emotion. *Cognition and emotion*, 17(2):297–314, 2003.
- Y. Lee, K. Son, T. S. Kim, J. Kim, J. J. Y. Chung, E. Adar, and J. Kim. One vs. many: Comprehending accurate information from multiple erroneous and inconsistent ai generations. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 2518–2531, 2024.
- Y. Rogers, W. R. Hazlewood, P. Marshall, N. Dalton, and S. Hertrich. Ambient influence: Can twinkly lights lure and abstract representations trigger behavioral change? In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, pages 261–270, 2010.
- M. Rudd, K. D. Vohs, and J. Aaker. Awe expands people’s perception of time, alters decision making, and enhances well-being. *Psychological science*, 23(10):1130–1136, 2012.
- M. N. Shiota, D. Keltner, and A. Mossman. The nature of awe: Elicitors, appraisals, and effects on self-concept. *Cognition and emotion*, 21(5):944–963, 2007.
- J. E. Stellar, A. M. Gordon, P. K. Piff, D. Cordaro, C. L. Anderson, Y. Bai, L. A. Maruskin, and D. Keltner. Self-transcendent emotions and their social functions: Compassion, gratitude, and awe bind us to others through prosociality. *Emotion Review*, 9(3):200–207, 2017.
- C. Zhang and B. Magerko. Generative ai literacy: A comprehensive framework for literacy and responsible use. *arXiv preprint arXiv:2504.19038*, 2025.
- C. Zhang, W. Wang, P. Pangaro, N. Martelaro, and D. Byrne. Generative image ai using design sketches as input: Opportunities and challenges. In *Proceedings of the 15th Conference on Creativity and Cognition*, pages 254–261, 2023.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our system—*Sound Clouds*, is an experimental art installation that probes the possibility of future possibilities of AML.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [NA]

Justification: This paper, as an art demo, does not have limitations.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This work has no theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: This paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: Paper does not include experiments requiring code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The art project conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes] .

Justification: We discussed how *Sound Clouds* act as a probe into humanistic-oriented approach to ambient intelligence.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We used a fine-tuned YOLO model, which we properly referenced.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [No]

Justification: We demoed the exhibit with the public, but due to the page limit, we did not give full detail about the participants’ interaction, but more focused on the details of the exhibit itself.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: The video recording, and the post-interaction user study is approved by our Institutional Review Board.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: We do not use LLM for anything in this research project except for formatting and grammarly check purpose.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.