USCILab3D: A Large-scale, Long-term, Semantically Annotated Outdoor Dataset

David S. Hippocampus[*]
Department of Computer Science
Cranberry-Lemon University
Pittsburgh, PA 15213
hippo@cs.cranberry-lemon.edu

Abstract

In this paper, we introduce the USC ILab 3D dataset, a comprehensive outdoor large-scale dataset designed for versatile applications across various domains, including computer vision, natural language processing, robotics, and machine learning. The dataset, USC ILab 3D, not only facilitates 3D reconstructions but also offers a diverse array of complex intersections for analysis. Despite covering a narrower geographical scope compared to Google Street View, our dataset prioritizes intricate intersections and boasts denser images and point clouds, enabling more precise 3D labeling and facilitating a broader spectrum of 3D vision tasks. Furthermore, we conduct benchmarking exercises on the USC ILab 3D dataset to evaluate the efficacy of current reinforcement learning and planning algorithms.

1 Introduction

With the recent advancements in 3D vision techniques, the integration of three-dimensional perception has become integral to various interdisciplinary domains. The progress in this field can be significantly propelled by leveraging large-scale datasets, which offer adaptability across a spectrum of downstream tasks. In this paper, we present USCILab3D—a large-scale, long-term, semantically annotated outdoor dataset of XX, labeled for XX categories. The excellence of our annotations is validated through Principal Preserved Component Analysis (PPCA). We aim to showcase the dataset's versatility by employing it in navigation experiments, demonstrating its efficacy across diverse domains.

2 Related work

Write something about the existing datasets

[*]Use footnote for providing further information about author (webpage, alternative address)—*not* for acknowledging funding agencies.

Submitted to the 37th Conference on Neural Information Processing Systems (NeurIPS 2023) Track on Datasets and Benchmarks. Do not distribute.

2.1 Comparison with similar datasets and features

2.2 Matterport

2.3 kitti semantic style

3 Dataset specification and collection.

3.1 Dataset collected over the entire USC campus

3.1.1 Data types.. RGB, Depth, pointcloud

3.1.2 Day timings

3.2 Perception rig specifics

3.3 Pose estimation and Panaroma stittching

d

3.4 Pairs of (dense 3d pointcloud, set of images) -> Semantically annotated (matterport3d, scannet style)

3.5 NeRF and Gaussian splatting because of groundtruth poses

3.6 Could be part of the training data for:

3.7 pairs of (scan, multi view rgb image)

Henghui Novel-view scene synthesis plays a crucial role in various applications, offering the potential for more realistic simulations. Traditional methods have inherent limitations that impede the creation of truly lifelike environments. In recent years, Radiance Field methods and Gaussian Splatting have emerged as a promising solution to address these limitations, significantly enhancing the quality of novel-view scene synthesis. We try to explore the capabilities of Radiance Field methods and Gaussian Splatting, and their potential to create a high-performance simulator that operates in real-time.

4 Method for semantic annotations for 3D pointclouds.

Recently using Radiance Field methods to synthesize novel-view scenes has a good performance on quality and speed. By synthesizing novel-view scenes, the agent can move in the scene like the real world instead of moving between images prepared before. It gives the agent higher degrees of freedom and more realistic simulations, which will help us reduce the gap between the simulator and the read-world. We want to create a simulator with extremely high performance, so rendering speed is one of the most important, we want it to be real-time.

5 Benchmarks

quality/speed

5.1 3D segmentation (semantic, panoptic, 4D panoptic, moving object, 3D scene completion)

5.2 NeRF and Gaussian splatting results

6 Conclusion

Acknowledgments and Disclosure of Funding

Use unnumbered first level headings for the acknowledgments. All acknowledgments go at the end of the paper before the list of references. Moreover, you are required to declare funding (financial activities supporting the submitted work) and competing interests (related financial activities outside the submitted work). More information about this disclosure can be found at: `https://neurips.cc/Conferences/2023/PaperInformation/FundingDisclosure`. You can use the `ack` environment provided in the style file. As opposed to the main NeurIPS track, acknowledgements do not need to be hidden.

References

References follow the acknowledgments. Use unnumbered first-level heading for the references. Any choice of citation style is acceptable as long as you are consistent. It is permissible to reduce the font size to `small` (9 point) when listing the references. Note that the Reference section does not count towards the page limit.

[1] Alexander, J.A. & Mozer, M.C. (1995) Template-based algorithms for connectionist rule extraction. In G. Tesauro, D.S. Touretzky and T.K. Leen (eds.), *Advances in Neural Information Processing Systems 7*, pp. 609–616. Cambridge, MA: MIT Press.

[2] Bower, J.M. & Beeman, D. (1995) *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SImulation System.* New York: TELOS/Springer–Verlag.

[3] Hasselmo, M.E., Schnell, E. & Barkai, E. (1995) Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *Journal of Neuroscience* **15**(7):5249-5262.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[TODO]**
   (b) Did you describe the limitations of your work? **[TODO]**
   (c) Did you discuss any potential negative societal impacts of your work? **[TODO]**

(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[TODO]**

2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? **[TODO]**
   (b) Did you include complete proofs of all theoretical results? **[TODO]**
3. If you ran experiments (e.g. for benchmarks)...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[TODO]**
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[TODO]**
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[TODO]**
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[TODO]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? **[TODO]**
   (b) Did you mention the license of the assets? **[TODO]**
   (c) Did you include any new assets either in the supplemental material or as a URL? **[TODO]**
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[TODO]**
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[TODO]**
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[TODO]**
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[TODO]**
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[TODO]**

A Appendix

Include extra information in the appendix. This section will often be part of the supplemental material. Please see the call on the NeurIPS website for links to additional guides on dataset publication.

1. Submission introducing new datasets must include the following in the supplementary materials:
   (a) Dataset documentation and intended uses. Recommended documentation frameworks include datasheets for datasets, dataset nutrition labels, data statements for NLP, and accountability frameworks.
   (b) URL to website/platform where the dataset/benchmark can be viewed and downloaded by the reviewers.
   (c) Author statement that they bear all responsibility in case of violation of rights, etc., and confirmation of the data license.
   (d) Hosting, licensing, and maintenance plan. The choice of hosting platform is yours, as long as you ensure access to the data (possibly through a curated interface) and will provide the necessary maintenance.
2. To ensure accessibility, the supplementary materials for datasets must include the following:

   (a) Links to access the dataset and its metadata. This can be hidden upon submission if the dataset is not yet publicly available but must be added in the camera-ready version. In select cases, e.g when the data can only be released at a later date, this can be added afterward. Simulation environments should link to (open source) code repositories.
   (b) The dataset itself should ideally use an open and widely used data format. Provide a detailed explanation on how the dataset can be read. For simulation environments, use existing frameworks or explain how they can be used.
   (c) Long-term preservation: It must be clear that the dataset will be available for a long time, either by uploading to a data repository or by explaining how the authors themselves will ensure this.
   (d) Explicit license: Authors must choose a license, ideally a CC license for datasets, or an open source license for code (e.g. RL environments).
   (e) Add structured metadata to a dataset's meta-data page using Web standards (like schema.org and DCAT): This allows it to be discovered and organized by anyone. If you use an existing data repository, this is often done automatically.
   (f) Highly recommended: a persistent dereferenceable identifier (e.g. a DOI minted by a data repository or a prefix on identifiers.org) for datasets, or a code repository (e.g. GitHub, GitLab,...) for code. If this is not possible or useful, please explain why.

3. For benchmarks, the supplementary materials must ensure that all results are easily reproducible. Where possible, use a reproducibility framework such as the ML reproducibility checklist, or otherwise guarantee that all results can be easily reproduced, i.e. all necessary datasets, code, and evaluation procedures must be accessible and documented.
4. For papers introducing best practices in creating or curating datasets and benchmarks, the above supplementary materials are not required.