

ReBound: An Open-Source 3D Bounding Box Annotation Tool for Active Learning

Wesley Chen[†], Andrew Edgley[†], Raunak Hota[†], Joshua Liu[†], Ezra Schwartz[†], Aminah Yizar[†], Neehar Peri[‡] and James Purtilo[‡]

Abstract

In recent years, supervised learning has become the dominant paradigm for training deep-learning based methods for 3D object detection. Lately, the academic community has studied 3D object detection in the context of autonomous vehicles (AVs) using publicly available datasets such as nuScenes and Argoverse 2.0. However, these datasets may have incomplete annotations, often only labeling a small subset of objects in a scene. Although commercial services exist for 3D bounding box annotation, these are often prohibitively expensive. To address these limitations, we propose ReBound, an open-source 3D visualization and dataset re-annotation tool that works across different datasets. In this paper, we detail the design of our tool and present survey results that highlight the usability of our software. Further, we show that ReBound is effective for exploratory data analysis and can facilitate active-learning. Our code and documentation is available on GitHub.

Keywords

Autonomous Driving, Active Learning, 3D Annotation Tools, Human Computer Interaction, Data Visualization

1. Introduction

3D object detection is a critical component of the autonomous vehicle (AV) perception stack [1, 2]. To facilitate research in 3D object detection, the AV industry has released large-scale 3D annotated multimodal datasets [2, 3, 4]. These datasets include LiDAR sweeps and multi-camera RGB images from diverse driving logs, which captures detailed information about the surrounding environment. Crucially, objects of interest are annotated by drawing 3D bounding boxes and labeling them as part of a particular category. Contemporary 3D detectors [5, 6, 7, 8, 9] are trained using supervised learning, and are limited by the annotations provided with the dataset. For example, objects in the nuScenes dataset are inconsistently labeled beyond 50m, making it challenging to evaluate long-range detection. Further, the nuScenes dataset does not label street signs and traffic lights, which are critical for safe navigation.

Manually re-annotating datasets with new 3D bounding box annotations is particularly challenging. Annotating 3D bounding boxes using 2D RGB images is difficult because it is not possible to accurately estimate bounding-box depth. Similarly, annotating 3D bounding boxes using LiDAR point clouds is difficult because LiDAR returns are sparse, making it tough to identify individual objects at long-range and in cluttered scenes, as shown in Fig. 2. Although commercial services can an-

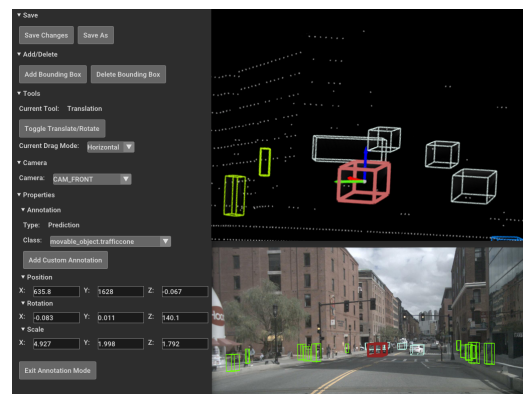


Figure 1: We present ReBound, an open-source 3D annotation tool that allows users to add, delete, and modify annotations from existing datasets or model predictions to support active learning.

notate 3D data at scale, it is often prohibitively expensive. As a result, several tools have been created for quick, efficient, and easy point cloud annotation [10, 11, 12, 13], with the intention of making it simpler for researchers to create their own datasets. However, these tools do not support a wide variety of data formats, multi-modal exploratory data analysis and active learning.

In this paper, we introduce ReBound, an open-source annotation tool which aims to simplify the process of multi-modal 3D visualization and bounding box annotation. Our framework is designed to provide users with the ability to import and modify annotations from various datasets. To achieve this, we propose a generic data type that can be extended to accommodate differ-

AutomationXP23: Intervening, Teaming, Delegating Creating Engaging Automation Experiences, April 23rd, Hamburg, Germany

*Corresponding author: purtilo@umd.edu

[†]These authors contributed equally.

[‡]These authors advised equally.

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)






Figure 2: Long-range 3D detection is crucial for safe navigation in autonomous vehicles. Modern benchmarks have greatly improved 3D object detection. However, standard benchmarks like nuScenes only evaluate detections up to 50m, masking the poor far-field object detection performance of 3D detectors (b & c). Existing benchmarks do not sufficiently annotate far-field objects (d & e), partly due to fewer LiDAR returns. Our proposed annotation framework facilitates far-field annotation, and will enable researchers to better study this problem. We adopt this figure from [14].

ent data formats. Additionally, ReBound enables active learning by allowing users to correct bounding box predictions and labels, create new custom labels, analyze model predictions, and export new annotations back to dataset-specific formats for model re-training. We measure the effectiveness of our tool through user studies, focusing on the ease of using the data conversion and annotation editing features. Our results show that the tool is both intuitive to use and useful for rapidly adding new annotations.

2. Related Work

In this section we present a brief overview of existing 3D annotation tools and active learning methods for object detection.

3D Annotation Tools. Modern 3D detectors require diverse, large-scale 3D annotations for supervised learning. A variety of tools have been created to address this challenge [15], including labelCloud[10], SAnE[11], LATTE[12], and 3D BAT[13], which all allow for efficient manual or semi-automatic dataset annotation. However, unlike these existing tools, ReBound also allows users to manually update annotations from existing datasets for evolving use cases like detecting a new category or updating annotations for far-field objects. Our tool currently supports re-annotation for the nuScenes[16], Waymo [17], and Argoverse 2.0[18] datasets.

Active Learning for 3D Object Detection. Although manually labeling tens of thousands of images or LiDAR sweeps can be prohibitively expensive, recent work in active learning suggests that we can bootstrap deep learning models by iteratively annotating and train-

ing on a targeted subset of informative examples to significantly improve model performance. Selecting the most relevant data to be annotated by human experts is a primary challenge for active learning. Recent works define a scoring function, using either an uncertainty-based [19, 20, 21, 22, 23, 24, 25, 26] or a diversity-based approach [27, 26], to determine the most informative training samples. When selecting samples based on uncertainty, informativeness is measured by the predictive uncertainty, and the samples with the highest uncertainty are provided to the human annotators. When selecting samples based on diversity, scores are assigned based on spatial and temporal diversity [27]. Both the uncertainty-based and diversity-based approaches have been used for bootstrapping 3D object detection, but the uncertainty-based approach has been shown to be more effective in practice. More recently, [26] attempts to combine both strategies. Despite substantial work in active learning for image recognition and 2D object detection, exploration into active learning for 3D object detection is limited. ReBound helps support active learning by allowing users to filtering predicted annotations based on detection confidence score, which is useful for measuring uncertainty.

3. ReBound: 3D Bounding-Box Re-Annotation

In this section, we describe the functionality and software architecture of ReBound. Our code, documentation, and videos demonstrating our tool are available on GitHub.

Data Format. ReBound currently supports three different datasets: nuScenes, Waymo Open Dataset, and

Argoverse 2.0. Annotations for these three datasets are converted into our generic ReBound format using a separate command line tool built in Python. Specifically, the generic data format captures the minimum information required to annotate a 3D bounding box (e.g. object center, box size, box rotation, and sensor extrinsics), which allows us to scale this format across multiple datasets. Users can support new datasets by extending our command line tool for their use case.

Data Visualization. The visualization tool has three windows: a control window, a point cloud viewer, and an RGB image viewer, as shown in Figure 1. The control window allows the user to navigate through different frames in a driving log, switch between different camera views in the RGB image window, and filter through both ground truth annotations and model predictions. The point cloud viewer displays the point cloud corresponding to the current frame, as well as all ground truth annotations and predictions (if available) for that frame as wireframe cuboids. The user can rotate, translate, and zoom in/out to interact with the scene. We use Open3D as our rendering back-end as this natively supports 3D rendering on top of LiDAR sweeps and RGB images.

Editing Tool. Users can directly click on a location in the point cloud window to add, edit, or delete an annotation. To edit a box’s properties, users must first click on a desired cuboid in the point cloud window. This will highlight the bounding box and display the position, rotation, size and annotation class of the selected box in the control window. These fields can all be directly updated, allowing the user to make precise changes. Users can also make coarse changes to the location of the selected box by clicking and dragging the box within the LiDAR viewer. To make it easier to interact with 3D objects, the mouse tool is restricted to two different control modes:

- Horizontal Translation Mode: The user can articulate objects along the X and Y axes with the Z axis locked
- Vertical Translation Mode: The user can click and drag objects along the Z axis and rotate objects about the Z axis.

Finer-grained modifications can be manually entered in the control window. Users can also delete the selected bounding box, and create new bounding boxes with a single mouse-click. All bounding box transformations are updated in both the LiDAR and RGB image viewers in real-time. In practice, the editing tool can be used to update bounding boxes that are too large, are misaligned with the true object location (c.f. Figure 4), update misclassified objects, and add new categories like stop signs and traffic lights.

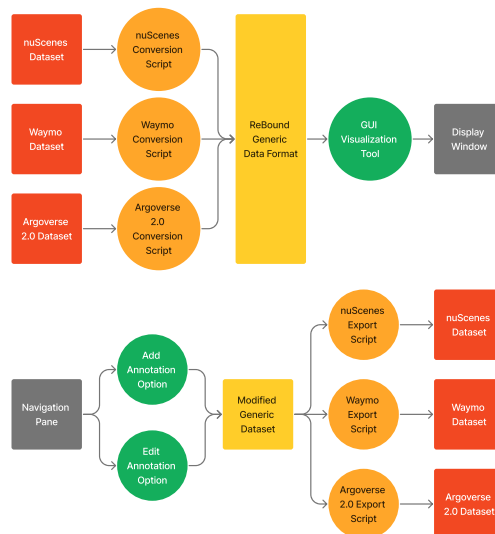


Figure 3: Our tool supports data conversion, visualization, and modification of the nuScenes, Waymo and Argoverse 2.0 datasets. First, dataset-specific RGB images, LiDAR sweeps, sensor extrinsics, and annotations are converted to the ReBound data format using our conversion scripts and visualized in the GUI. Using this GUI, we can add, edit, or delete existing annotations. Lastly, we can export data from the ReBound data format back to the respective dataset-specific formats using the provided export scripts.

Exporting. Users can export updated annotations back to the original dataset-specific formats using a command line tool. Importantly, ReBound only exports the modified bounding box annotations back to the original format. Since ReBound does not modify the LiDAR sweeps or RGB images, we can directly use the data from the original dataset. We find that this dramatically increases data export speed. Interestingly, since we convert all datasets into a unified format, we can easily export annotations *between* datasets. Concretely, this means we can convert nuScenes annotations into the Argoverse 2.0 format using the ReBound format as an intermediary, making it easier to evaluate models across different datasets. Further, this can facilitate future research in model robustness and domain adaptation between different datasets. Similar to the data conversion tool that converts from dataset-specific formats to the ReBound format, users that wish to support a new dataset can extend our command line tool for their use case.

4. Experiments

In this section we present the results of our user survey to evaluate the effectiveness of our tool.

Ease of Use. We conducted surveys to evaluate the tool’s ease of use in facilitating data conversion between data formats and bounding box adjustments among both people familiar and unfamiliar with autonomous vehicle datasets and 3D annotation tools.

We asked ten participants to perform 13 tasks (shown in Table 1) over a video conference. Participants were shown a demonstration of the tool before being asked to complete a new, but related task using the tool. This method of demonstration allows us to avoid re-downloading the datasets and re-installing the software on a new computer for each trial. After completing all 13 tasks, users were asked to complete a questionnaire designed to rate different parts of the user experience on a scale of one to five (five is highest).

Based on our survey, we identified that the greatest challenge for users was translating and rotating bounding boxes. Both of these operations require fine-grained controls and navigating in a 3D space, which may be not be intuitive for some users. However these features also had high standard deviation compared to the other tasks, indicating that the utility of our tool may depend on prior use of 3D visualization and annotation tools. In general, participants stated that the application can be useful for autonomous vehicle research.

Bounding Box Adjustments. We visualize ground truth annotations in both the LiDAR and RGB image viewers. We are able to identify examples where the ground truth annotation is misaligned with the real object, as shown in Figure 4, highlighting a potential application of our tool. This is practically meaningful as training 3D detectors with noisy labels can result in degraded model performance.

Task List	Avg.	Std. Dev
Convert nuScenes data to the ReBound generic type	4.1	0.74
Enter annotation mode	4.2	0.79
Add a new annotation bounding box	3.8	1.4
Add a custom annotation type	4.7	0.67
Select an existing bounding box	4.5	0.97
Translate a bounding box	3.8	1.03
Rotate a bounding box	3.6	1.17
Change the label of a bounding box	4.8	0.42
Delete a bounding box	4.8	0.63
Use the control viewer to edit an annotation	4	1.15
Save modified annotations	4.5	0.71
Exit the application	4.5	0.53
Export data back to the nuScenes data format	3.9	0.88

Table 1

Our survey results show that users found it easier to create, delete, and modify annotations but found it more difficult to rotate, translate, and export bounding boxes from the ReBound data format to dataset-specific formats.

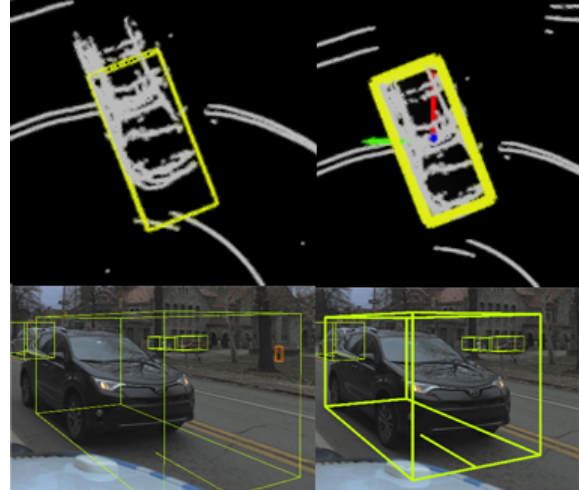


Figure 4: We visualize a ground truth annotation from the Argoverse 2.0 dataset (left). We find that the annotation is misaligned with the true car location. We adjust this ground truth annotation using ReBound (right). Visually inspecting the updated bounding box in both the LiDAR and RGB image viewer confirms that the updated annotation correctly localizes the car.

5. Conclusion

The academic community studies 3D object detection using publicly available AV datasets, which often provide LiDAR sweeps, RGB images, and 3D bounding box annotations. However, these datasets are limited by the annotations provided by their curators. In some cases, we find that these annotations may be incorrect or incomplete. Existing annotation tools cannot be easily used to update annotations for existing datasets, which makes it difficult to fix incorrect annotations, annotate new categories of interest, or iteratively improve 3D object detection models through active learning. In this paper, we propose ReBound, an open-source alternative that simplifies the process of bounding box re-annotation for existing datasets. We validate the utility of our tool through user surveys and find that users are able to rapidly add, modify, and delete 3D bounding box annotations.

Acknowledgments

Author James Purtilo was supported by N000142112821 while working on this project. This work is supported by the SEAM Lab at the University of Maryland and the Argo AI Center for Autonomous Vehicle Research at Carnegie Mellon University.

References

- [1] A. Geiger, P. Lenz, R. Urtasun, Are we ready for autonomous driving? the kitti vision benchmark suite, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [2] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, O. Beijbom, nuscenes: A multimodal dataset for autonomous driving, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [3] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, J. Hays, Argoverse 2: Next generation datasets for self-driving perception and forecasting, in: *Neural Information Processing Systems Datasets and Benchmarks Track*, 2021.
- [4] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, D. Anguelov, Scalability in perception for autonomous driving: Waymo open dataset, in: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [5] T. Yin, X. Zhou, P. Krähenbühl, Center-based 3d object detection and tracking, *arXiv preprint arXiv:2006.11275* (2020).
- [6] N. Peri, A. Dave, D. Ramanan, S. Kong, Towards long-tailed 3d detection, in: *Conference on Robot Learning (CoRL)*, 2022.
- [7] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, O. Beijbom, Pointpillars: Fast encoders for object detection from point clouds, in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [8] B. Zhu, Z. Jiang, X. Zhou, Z. Li, G. Yu, Class-balanced grouping and sampling for point cloud 3d object detection, *arXiv preprint arXiv:1908.09492* (2019).
- [9] N. Peri, J. Luiten, M. Li, A. Osep, L. Leal-Taixe, D. Ramanan, Forecasting from lidar via future object detection, *arXiv:2203.16297* (2022).
- [10] C. Sager, P. Zschech, N. Köhl, labelcloud: A lightweight domain-independent labeling tool for 3d object detection in point clouds, *CoRR abs/2103.04970* (2021).
- [11] H. A. Arief, M. Arief, G. Zhang, Z. Liu, M. Bhat, U. G. Indahl, H. Tveite, D. Zhao, Sane: Smart annotation and evaluation tools for point cloud data, *IEEE Access* 8 (2020) 131848–131858.
- [12] B. Wang, V. Wu, B. Wu, K. Keutzer, LATTE: accelerating lidar point cloud annotation via sensor fusion, one-click annotation, and tracking, in: *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019, Auckland, New Zealand, October 27-30, 2019, IEEE*, 2019, pp. 265–272.
- [13] W. Zimmer, A. Rangesh, M. M. Trivedi, 3d BAT: A semi-automatic, web-based 3d annotation toolbox for full-surround, multi-modal data streams, in: *2019 IEEE Intelligent Vehicles Symposium, IV 2019, Paris, France, June 9-12, 2019, IEEE*, 2019, pp. 1816–1821.
- [14] S. Gupta, J. Kanjani, M. Li, F. Ferroni, J. Hays, D. Ramanan, S. Kong, Far3det: Towards far-field 3d detection, in: *NeurIPS*, 2022.
- [15] Y. Li, J. Ibañez-Guzmán, Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems, *IEEE Signal Process. Mag.* 37 (2020) 50–61.
- [16] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, O. Beijbom, nuscenes: A multimodal dataset for autonomous driving, in: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, Computer Vision Foundation / IEEE*, 2020, pp. 11618–11628.
- [17] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, D. Anguelov, Scalability in perception for autonomous driving: Waymo open dataset, in: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, Computer Vision Foundation / IEEE*, 2020, pp. 2443–2451.
- [18] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, J. Hays, Argoverse 2: Next generation datasets for self-driving perception and forecasting, in: J. Vanschoren, S. Yeung (Eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021.
- [19] E. Hausmann, M. Fenzi, K. Chitta, J. Ivanec, H. Xu, D. Roy, A. Mittel, N. Koumchatzky, C. Farabet, J. M. Alvarez, Scalable active learning for object detection, in: *IEEE Intelligent Vehicles Symposium, IV 2020, Las Vegas, NV, USA, October 19 - November 13, 2020, IEEE*, 2020, pp. 1430–1435.
- [20] T. Yuan, F. Wan, M. Fu, J. Liu, S. Xu, X. Ji, Q. Ye, Multiple instance active learning for object detection, in: *CVPR*, 2021.
- [21] D. Feng, X. Wei, L. Rosenbaum, A. Maki, K. Dietmayer, Deep active learning for efficient training of a lidar 3d object detector, in: *2019 IEEE Intelligent Vehicles Symposium, IV 2019, Paris, France, June*

- 9-12, 2019, IEEE, 2019, pp. 667–674.
- [22] S. Roy, A. Unmesh, V. P. Namboodiri, Deep active learning for object detection, in: British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018, BMVA Press, 2018, p. 91.
 - [23] A. Hekimoglu, M. Schmidt, A. Marcos-Ramiro, G. Rigoll, Efficient active learning strategies for monocular 3d object detection, in: 2022 IEEE Intelligent Vehicles Symposium, IV 2022, Aachen, Germany, June 4-9, 2022, IEEE, 2022, pp. 295–302.
 - [24] S. Schmidt, Q. Rao, J. Tatsch, A. C. Knoll, Advanced active learning strategies for object detection, in: IEEE Intelligent Vehicles Symposium, IV 2020, Las Vegas, NV, USA, October 19 - November 13, 2020, IEEE, 2020, pp. 871–876.
 - [25] M. Meyer, G. Kusch, Automotive radar dataset for deep learning based 3d object detection, in: 2019 16th European Radar Conference (EuRAD), 2019.
 - [26] Y. Luo, Z. Chen, Z. Wang, X. Yu, Z. Huang, M. Bakdashmotlagh, Exploring active 3d object detection from a generalization perspective, CoRR abs/2301.09249 (2023).
 - [27] Z. Liang, X. Xu, S. Deng, L. Cai, T. Jiang, K. Jia, Exploring diversity-based active learning for 3d object detection in autonomous driving, CoRR abs/2205.07708 (2022).