

TBD Pedestrian Data Collection: Towards Rich, Portable, and Large-Scale Natural Pedestrian Data

Allan Wang¹, Daisuke Sato¹, Yasser Corzo¹, Sonya Simkin¹, Abhijat Biswas¹, and Aaron Steinfeld¹

Abstract—Social navigation and pedestrian behavior research has shifted towards machine learning-based methods and converged on the topic of modeling inter-pedestrian interactions and pedestrian-robot interactions. For this, large-scale datasets that contain rich information are needed. We describe a portable data collection system, coupled with a semi-autonomous labeling pipeline. As part of the pipeline, we designed a label correction web application that facilitates human verification of automated pedestrian tracking outcomes. Our system enables large-scale data collection in diverse environments and fast trajectory label production. Compared with existing pedestrian data collection methods, our system contains three components: a combination of top-down and ego-centric views, natural human behavior in the presence of a socially appropriate “robot”, and human-verified labels grounded in the metric space. To the best of our knowledge, no prior data collection system has a combination of all three components. We further introduce our ever-expanding dataset from the ongoing data collection effort – the *TBD Pedestrian Dataset* and show that our collected data is larger in scale, contains richer information when compared to prior datasets with human-verified labels, and supports new research opportunities.

I. INTRODUCTION

Pedestrian datasets are essential tools for modeling socially appropriate robot behaviors, recognizing and predicting human actions, and studying pedestrian behavior. Researchers may use these data to predict future pedestrian motions, including forecasting their trajectories [1], [2], [3], and/or navigation goals [4], [5]. In social navigation, datasets can also be used to model [6], [7] or evaluate robot navigation behavior [8]. For this, an in-the-wild pedestrian dataset that is large-scale and supports ground-truth metric labels is desired.

However, existing public pedestrian datasets are either unlabelled [9], [10], only rely on labels produced by an automated pipeline [11], [12], only contain pixel level information [13], [14] or are small in scale [15], [16], [17]. We propose a system that can collect large quantities of quality data efficiently. The data collected using our system feature a novel full combination of three critical elements: a combination of top-down and ego-centric views, natural human motion, and human-verified labels grounded in the metric space. This allows the data collected using our system to contain rich information.

Large datasets with quality labels and rich information can assist in addressing human behavioral research questions

¹The authors are with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, USA {allanwan, daisukes, ycorzo, ssimkin, abhijatb}@andrew.cmu.edu, steinfeld@cmu.edu




Fig. 1: This set of images represent the same moment recorded from multiple sensors: a) Top-down view image taken by a static camera with grounded pedestrian trajectory labels shown. b) Ego-centric point cloud from a 3D LiDAR with the projected trajectories from (a). c) Ego-centric RGB and depth images from the mounted stereo camera. Green vertical bars represent the projected labels. Note that two pedestrians at the back are partially and completely occluded from the stereo camera.

that require the modeling of interaction. For example, a key problem researchers have tried to address is the *freezing robot problem* [18]. Researchers have attributed this problem to the robot’s inability to model interactions [19]. Works such as [20] have used datasets to show that modeling the anticipation of human reactions to the robot’s actions enables the robot to deliver a better performance. However, interactions are diverse and uncommon in human crowds, they contain many types [21] (e.g. passing, following, grouping) and can further be diversified by the environment (e.g. an open plaza or a narrow corridor), so pedestrian datasets need to be large-scale in order to capture enough interaction data.

Autonomous vehicle datasets [22], [23] have inspired a plethora of research. However, a dataset of similar caliber and label quality in pedestrian-dominant environments has yet to arrive. As a step toward this goal, we have constructed a data collection system that can achieve these two requirements: large data quantity and diversity, and human-verified positional labels. First, we ensure that our equipment is portable and easy to set up. This allows collecting data in a variety of locations with limited lead time. Second, we address the challenge of labeling large quantities of data using a semi-autonomous labeling pipeline. We employ

a state-of-the-art deep learning-based [24] tracking module combined with a human inspection and tracking error-fixing web application to semi-automatically produce high-quality ground truth pedestrian trajectories in metric space. We make the web application open-source¹ so that other researchers can use this tool or contribute to this effort.

While we hope our contributions support robot system improvements in the community and we aim to accommodate a wide variety of pedestrian behavior research, our dataset primarily supports human environment navigation research that requires ground truth pedestrian positional information, such as social navigation, pedestrian trajectory prediction, and ego-centric perception. Specifically, we include three important characteristics. (1) Top-down view and ego-centric views: This ensures that the robot has access to ground-truth data even with occlusions. (2) Natural human motion: The manual pushing of the inconspicuous suitcase robot mitigates the curiosity effects of nearby pedestrians. And (3) Ground truth labeling in metric space: This allows our dataset to be useful for research where positional pedestrian data are needed. To the best of our knowledge, publicly available datasets only have at most two of these characteristics.

We demonstrate our system through a dataset collected in a large indoor space: the TBD Pedestrian Dataset². Our dataset contains scenes with a variety of crowd densities and pedestrian interactions. We show through our analysis that our dataset (Batch 1: 133 minutes - 1416 trajectories; Batch 2: 626 minutes - 10300 trajectories) is larger in scale and contains unique characteristics compared to prior similar datasets. This is an ongoing effort and we plan to collect additional data in more diverse locations.

II. RELATED WORK

With the explosion of data-hungry machine learning methods in robotics, demand for pedestrian datasets has been on the rise in recent years. One popular category of research in this domain is human trajectory prediction (e.g., [1], [2], [3], [4], [5], [25], [26], [27]). Much of this research utilizes selected mechanisms to model pedestrian interactions in hopes for better prediction performance (e.g., pooling layers in the deep learning frameworks [1], [2] or graph-based representations [26]). Rudenko et al. [28] provides a good summary into this topic. While the state-of-the-art performance keeps improving with the constant appearance of newer models, it is often unclear how well these models can generalize in diverse environments. As shown in [28], many of these models only conduct their evaluation on the relatively small-scale ETH [15] and UCY [16] datasets.

Another popular demand for pedestrian datasets comes from social navigation research. Compared to human motion prediction research, social navigation research focuses more on planning. For example, social navigation research uses learning-based methods to identify socially appropriate motion for better robot behavior, such as deep reinforcement




Fig. 2: Sensor setup used to collect the TBD Pedestrian Dataset. (left) One of the nodes used to capture top-down RGB views. (middle) The cart used to capture ego-centric sensor views during data collection for Set 1. (right) The suitcase robot used to capture ego-centric sensor views during data collection for Set 2.

learning [29], [30], [31] or inverse reinforcement learning [6], [32]. Due to the lack of sufficiently large datasets, these models often train in simulators that lack realistic pedestrian behavior. Apart from training, datasets are also increasing in popularity in social navigation evaluation due to their realistic pedestrian behavior [33]. Social navigation methods are often evaluated in environments using pedestrian data trajectory playback (e.g., [19], [34], [35], [36]). However, similar to human motion prediction research, these evaluations are typically only conducted on the ETH [15] and UCY [16] datasets, as shown by [33]. These two datasets only use overhead views, and therefore lack the ego-centric view used by most robots.

Large-scale and high-quality datasets exist for other navigation-related applications and research. Autonomous vehicle datasets such as nuScenes [22], Cityscapes [23] and ArgoVerse [37] also contain pedestrian-related data. However, pedestrians often have limited appearances on sidewalks or at crosswalks. There is also no data on how pedestrians navigate indoors. Another group of similar datasets mainly supports computer vision-related research, such as MOT [38] for pedestrian tracking, and Stanford Drone Dataset (SDD) [13] and VIRAT [14] for pedestrian motion/goal prediction on the image level. Detailed comparisons of the characteristics between the TBD Pedestrian Dataset and similar existing datasets can be found in section IV-A.

Simulators can fill in the role of datasets for both training and evaluation. Simulators such as SocNavBench [8], CrowdNav [30], PedSIM [39], and SEAN [40] are in use by the research community. However, sim-to-real transfer is an unsolved problem in robotics. Apart from lack of fidelity in visuals and physics, pedestrian simulators in particular entail the additional paradox of pedestrian behavior realism [41]: If pedestrian models are realistic enough for use in simulators, why don't we apply the same model to social navigation?

III. SYSTEM DESCRIPTION

A. Hardware Setup

Our system supports multiple static FLIR Blackfly RGB cameras for labeling and metric space calculations (Figure 2). The TBD Pedestrian Dataset contains two different configuration sets for the same physical space. Both sets include three cameras surrounding the scene on the upper floors

¹https://github.com/CMU-TBD/tbd_label_correction_UI

²<https://tbd.ri.cmu.edu/tbd-social-navigation-datasets>

overlooking the ground level at roughly 90 degrees apart from each other (Figure 3). The RGB cameras are connected to portable computers powered by lead-acid batteries. The multiple cameras complement each other’s coverage, because even from overhead views, partial occlusions occur. We used the three overhead view cameras to label trajectories. We also positioned three more units on the ground floor, but did not use them for pedestrian labeling. The cameras’ intrinsics are precalibrated and their extrinsics are obtained via physical landmarks that we placed in the scene.

In addition to static cameras, we pushed a cart (Set 1) or robotic suitcase [42] (Set 2) through the scene. The cart (Figure 2) was equipped with a ZED stereo camera to collect ego-centric RGB-D views of the scene. A GoPro Fusion 360 camera for capturing high definition 360 videos of nearby pedestrians was mounted above the ZED. Data from the 360 camera is useful for capturing pedestrian pose data and facial expressions for future work. The ZED camera was powered by a laptop with a power bank. Our entire data collection hardware system is portable and does not require power outlets, thereby allowing data collection outdoors or in areas where wall power is inaccessible.

The robotic suitcase (Set 2) is a converted carry-on rolling suitcase. It is equipped with an IMU and a Velodyne VLP-16 3D LiDAR sensor. In addition, the same ZED camera and 360 camera are mounted on the suitcase handle. The robot’s computer, batteries, and all its internal components are hidden inside the suitcase, so pushing the robot resembles pushing a suitcase. We selected this robot because of its inconspicuous design to reduce curious, unnatural reactions from nearby pedestrians, as curious pedestrians may intentionally block robots or display other unnatural movements [43]. While it is true that real-world pedestrians will react to mobile robots curiously in the short term and some may argue in favor of a more robotic appearance, we envision that such curiosity effects will die down in the long term.

During certain data collection sessions, we pushed the cart or the suitcase robot from one end of the scene to another end, while avoiding pedestrians and obstacles along the way in a natural motion similar to a human pushing a delivery cart or walking with a suitcase. This collects ego-centric views from a mobile robot traversing through the human environment. However, unlike other datasets such as [9], [10], [17], and [44] that use a tele-operated robot, or [45] that uses a scripted policy to act autonomously, we chose to have all motion performed by the human walking with the system. This provides better trajectory control, increased safety, and further reduced the novelty effect from surrounding pedestrians.

Both sets of data collection occurred on the ground level in a large indoor atrium area (Figure 3). Half of the atrium area has fixed entry/exit points that lead to corridors, elevators, stairs, and doors to the outside. The other half of the atrium is adjacent to another large open area and is unstructured with no fixed entry/exit points. We collected data around lunch and dinner times to ensure higher crowd densities. Additional data collection are planned at more diverse locations.



Fig. 3: Hardware setup for the TBD Pedestrian Dataset. Blue circles indicate positions of RGB cameras. Green box shows our suitcase robot pushed through the scene. The white area is where trajectory labels are collected.

B. Post-processing Pipelines

To ensure time synchronization across the captured videos, we employed Precision Time Protocol over a wireless network to synchronize each of the computers powering the cameras. In addition, we used an LED light and check for the light signal during the post-processing stage to synchronize the frames of all captured data for each recording session.

The next step is to localize the cart/suitcase in the scene. For Set 1, this was achieved by identifying the cart/suitcase on the overhead camera videos and then applying the camera matrices to obtain the metric coordinates. For Set 2, we first made a map inside the building and then computed its location in the post processing phase by utilizing the robotic suitcase software³ powered by Cartographer⁴.

For pedestrian tracking, we first tracked the pedestrians on the overhead camera videos. We found ByteTrack [24] to be very successful in tracking pedestrians in the image space. Upon human verification over our entire data, ByteTrack successfully tracks 95.1% of the pedestrians automatically.

For the automatically tracked labels in pixel space, we needed to convert them into metric space. Each camera video contained a set of tracked trajectories in the image space. We estimated the 3D trajectory coordinates for each pair of 2D trajectories from different cameras and the set of estimated coordinates that resulted in the lowest reprojection error were selected to be the trajectory coordinates in the metric space.

C. Human Label Verification

To ensure the label quality of the data, human verification of the tracked trajectories from ByteTrack is desired. Semi-autonomous labeling procedures are common in autonomous driving datasets and pedestrian datasets. However, surveying existing pedestrian datasets literature, we noticed that datasets that do contain human-verified metric space labels are often relatively small [15], [16], [17], [46], and large-scale datasets often only use automated tracking pipelines [11], [12], [14] or do not label surrounding pedestrians [9],

³<https://github.com/cmu-cabot>

⁴<https://github.com/cartographer-project/cartographer>



Fig. 4: Application interface for the human verification process. It contains a media player and various options to fix tracking errors automatically and manually.

[10]. We attribute this to a lack of tools to streamline the human verification process.

To this end, we designed an open-source web application (Figure 4) using Matlab App Designer. The tool was designed to minimize complete human relabeling of erroneously tracked trajectories. The application contains a media player. When using the application, human labelers only need to watch videos with the automatically tracked trajectories. When an error is noticed, the labeler only needs to indicate to the system the type and location of the error. The system then fixes the errors in the background, and updates the trajectory visualization accordingly. Currently, the application contains the following set of error-fixing options:

- **Break:** Used when ByteTrack incorrectly assigns the same trajectory to two different pedestrians.
- **Join:** Used when two different trajectories actually belong to the same pedestrian.
- **Delete:** Used when a ghost trajectory appears, such as incorrectly tracking an unworn jacket as a pedestrian.
- **Disentangle:** Used when the two trajectories of two pedestrians swapped in the middle, which can happen when one partially occludes the other.

The web application also supports undoing previous actions, partial or complete relabeling of trajectories, and labeling missing trajectories. For future work, we are looking at possible platforms to launch the application so that the human verification process can be a crowdsourcing effort.

Combined with ByteTrack, it took an expert labeler about 30 hours to produce human-verified labels for 375K frames of data, or 10300 trajectories. ByteTrack successfully tracks 95.1% of trajectories. Among the trajectory errors that require human rectification, 0.35% are fixed by “Break”, 1.25% are fixed by “Join”, 0.29% are fixed by “Delete”, 0.89% are fixed by “Disentangle”, 1.21% are fixed by “Relabel”, and 0.48% are fixed by “Missing”.

TABLE I: A survey of existing pedestrian datasets on how they incorporate the three components in section IV-A. For component 1, a “No” means either not human verified or not grounded in metric space. For component 2, TD stands for “top-down view” and “E” stands for “ego-centric view”.

Datasets	Comp. 1 (metric labels)	Comp. 2 (views)	Comp. 3 (“robot”)
TBD (Ours)	Yes	TD + E	Human + Robot
ETH [15]	Yes	TD	N/A
UCY [16]	Yes	TD	N/A
Edinburgh Forum [12]	No	TD	N/A
VIRAT [14]	No	TD	N/A
Town Centre [47]	No	TD	N/A
Grand Central [48]	No	TD	N/A
CFF [49]	No	TD	N/A
Stanford Drone [13]	No	TD	N/A
L-CAS [44]	No*	E	Robot
WildTrack [46]	Yes	TD	N/A
JackRabbit [17]	Yes	E	Robot
ATC [11]	No	TD	N/A
THÖR [45]	Yes	TD + E	Robot
SCAND [9]	No	E	Robot
Crowd-Bot [10]	No	E	Human + Robot

IV. DATASET CHARACTERISTICS AND ANALYSIS

A. Comparison with Existing Datasets

Compared to existing datasets collected in natural pedestrian-dominant environments, our TBD pedestrian dataset contains three components that greatly enhance the dataset’s utility. These components are:

Human verified labels grounded in metric space. ETH [15] and UCY [16] datasets are the most popular datasets among human behavior analysis papers [28]. We believe this is partly because the trajectory labels in these datasets are human verified and are grounded in metric space rather than pixel space (e.g. [13] and [47] only contain labels in bounding boxes). Having labels grounded in metric space eliminates the possibility that camera poses might have an effect on the scale of the labels. It also makes the dataset useful for robot navigation related research because robots plan in the metric space rather than pixel space.

Combination of top-down views and ego-centric views. Similar to datasets with top-down views, we used top-down views to obtain ground truth trajectory labels for every pedestrian present in the scene. Similar to datasets with ego-centric views, we gathered ego-centric views from a “robot” to imitate robot perception of human crowds. A dataset that contains both top-down views and ego-centric views will be useful for research projects that rely on ego-centric views. This allows ego-centric inputs to their models, while still having access to ground truth knowledge of the entire scene.

Naturalistic human behavior with the presence of a “robot”. Unlike datasets such as [9], [10], [17], and [44], the “robot” that provides ego-centric view data collection is a cart or a suitcase robot being pushed by human. As mentioned in section III-A, doing so reduces the novelty effects from the surrounding pedestrians. Having the “robot” pushed by humans also ensures safety for the pedestrians and its own motion has less jerk and more humanlike behavior.

TABLE II: Dataset comparison statistics, for those with human verified labels grounded in metric space. Numbers in parenthesis are for data that includes the ego-centric view.

Datasets	Duration	# Trajectories	Label Freq (Hz)
TBD Set 1	133 (51) min	1416	60
TBD Set 2	626 (213) min	10300	10
ETH [15]	25 min	650	15
UCY [16]	16.5 min	786	2.5
WildTrack [46]	200 sec	313	2
JackRabbit [17]	62 min	260	7.5
THÖR [45]	60+ min	600+	100

TABLE III: Comparison of statistics between our dataset and other datasets according to the methods in [45].

Datasets	Tracking Duration [s]	Percep. Noise [ms^{-2}]	Motion Speed [ms^{-1}]	Min Dist. To Ped. [m]
TBD Set 2	25.6 ± 57.1	0.55	0.88 ± 0.52	1.25 ± 1.44
THÖR [45]	16.7 ± 14.9	0.12	0.81 ± 0.49	1.54 ± 1.60
ETH [15]	9.4 ± 5.4	0.19	1.38 ± 0.46	1.33 ± 1.39
ATC [11]	39.7 ± 64.7	0.48	1.04 ± 0.46	0.61 ± 0.16
Edinburgh [12]	10.1 ± 12.7	0.81	1.0 ± 0.64	3.97 ± 3.5

As shown in Table I, current datasets only contain at most two of the three components⁵. A close comparison is the THÖR dataset [45], but its ego-centric view data are collected by a robot running on predefined trajectories. Additionally, unlike all other datasets in Table I, the THÖR dataset is collected in a controlled lab setting rather than in the wild. This injects artificial factors into human behavior.

B. Dataset Size

Table II demonstrates the ability of a semi-automatic labeling pipeline to produce large amounts of data. With the aid of an autonomous tracker, humans only need to verify and make occasional corrections on the tracking outcomes instead of locating the pedestrians on every single frame. The data we have collected so far surpassed all other datasets that provide human-verified labels in the metric space in terms of total time and number of pedestrians. We will continue this effort and collect more data for future work.

C. Dataset Statistics

Extending the evaluations performed in THÖR [45], we added the same suite of analysis on Set 2 of our TBD dataset. The evaluation metrics were the following. (1) *Tracking Duration (s)*: Average time duration of tracked trajectories. (2) *Perception Noise (ms^{-2})*: The average absolute acceleration of the trajectories. (3) *Motion Speed (ms^{-1})*: Velocities of the trajectories measured in 1 second intervals. (4) *Minimum Distance Between People (m)*: Minimum Euclidean distance between two closest observed people.

As shown in Table III, our dataset has considerable average trajectory duration (± 25.6) and large variation (± 57.1),

⁵*L-CAS dataset does provide human verified labels grounded in the metric space. However, its pedestrian labels do not contain trajectory data.

TABLE IV: Trajectory prediction displacement error on ETH/UCY datasets and TBD dataset Set 2.

Models	ETH/UCY Dataset			
	Static + Dynamic		Dynamic	
	ADE(m)	FDE(m)	ADE(m)	FDE(m)
Social-GAN [2]	0.48	0.96	0.59	1.13
Trajectron++ [50]	0.27	0.49	0.35	0.65
AgentFormer [51]	0.23	0.39	0.25	0.44
TBD Set 2				
Social-GAN	0.36	0.72	0.64	1.30
Trajectron++	0.16	0.28	0.43	0.83
AgentFormer	0.15	0.23	0.30	0.52

second only to ATC, which has a $900m^2$ coverage. While our dataset has a much smaller coverage, we attribute this to the presence of pedestrians changing navigation goals and static pedestrians in our dataset. Static pedestrians include standing pedestrians having conversations or pedestrians sitting on chairs. Their presence in our dataset often has a long duration, which also causes big variation in this metric. The tracking noise of our system was sub-optimal when compared to ETH and THÖR, which is likely due to noisy tracking of the sitting pedestrians. We observed that sitting pedestrians change their body poses frequently, which causes the tracked bounding boxes to change size frequently. We will investigate how to improve this for future work. The motion speeds of our dataset trajectories are lower (0.88), which suggests the presence of more static pedestrians. We also have the second-highest variation in motion speed (± 0.52), suggesting that our dataset captures a wide range of pedestrian behavior. From the minimum distance between people, it can be inferred that our dataset captures both dense and sparse population scenarios, as indicated by the middle mean value (1.25) among the others and high variance (± 1.44). Note that [45] also measures trajectory curvatures, but we noticed that this measurement is heavily affected by how static pedestrians are processed. [45] does not provide any details on this, so we decided not to evaluate this metric.

D. Behavior Distribution Analysis

We additionally leveraged trajectory prediction models to evaluate our dataset. Trajectory prediction models' performance has advanced significantly. We believe these well-trained models can be utilized in the other ways, such as characterizing the variety of pedestrian behavior in datasets. Almost all trajectory prediction models have been tuned and trained on ETH/UCY datasets. Some have additionally made predictions on SDD [13] or autonomous driving datasets. Because we were primarily concerned with metric labels and pedestrian environments, we did not evaluate models trained on SDD or autonomous vehicle datasets. We also chose models that largely leverage pedestrian positional data and can work independently without image patch inputs [25] or semantic segmentation [52].

Our dataset contains simple sessions and challenging sessions. To test if our dataset contains pedestrian behavior outside of ETH/UCY dataset domains, we analyzed the




Fig. 5: Examples from the TBD Set 1. a) a dynamic group. b) a static conversational group. c) a tour group with 14 pedestrians. d) a pedestrian affecting other pedestrians by asking them to come to the table. e) pedestrians stop and look at their phones. f) two pedestrians change their navigation goals and turn toward the table. g) a group of pedestrians change their navigation goals multiple times. h) a crowded scene where pedestrians are heading in different directions.

sessions from our dataset with great variety in pedestrian behavior for this evaluation. We selected Social-GAN [2] as the baseline model and Trajectron++ [50] and AgentFormer [51] as relatively state-of-the-art models. Because the models trained on each of the other four sub-datasets did not perform significantly differently in our dataset, we only report the average *Average Displacement Error* (ADE) and the average *Final Displacement Error* (FDE) across the five models.

We observed that when all pedestrians are concerned, the prediction models all perform better in our dataset (Table IV). We believe it can be attributed to the larger presence of static pedestrians in our datasets compared to ETH/UCY because it is unlikely to yield large errors when predicting future trajectories of static pedestrians. We additionally defined dynamic pedestrians as pedestrians who move at least $1m$ during the prediction window. We included static pedestrians during model inference but only dynamic pedestrians were considered for evaluation. With this, we discovered that all the prediction models’ performances degrade. This indicates that the models have encountered more unseen scenarios in our dataset and that the moving pedestrians in our dataset exhibit more diverse navigation behavior and wider behavior distribution compared to the ones in ETH/UCY.

E. Qualitative Pedestrian Behavior

Due to the nature of the environment where we collect the data, we observe a mixture of corridor and open-space pedestrian behavior, many of which are rarely seen in other datasets. As shown in Figure 5, we observe both static conversation groups and dynamic walking groups. We also observe that some pedestrians change goals mid-navigation. During certain sessions, some pedestrians set up activity areas with tables and chairs to engage with passing pedestrians. This creates more interesting interaction scenarios.

V. CONCLUSION

This paper presents a data collection system that is portable and enables large-scale data collection. This paper

also presents a human label verification tool that streamlines the labeling process. Our semi-autonomous pipeline easily produces human-verified labels in order to meet the demands of the large-scale data collected by our hardware. Our systems offer better utility for pedestrian behavior research because our systems consist of human-verified labels grounded in the metric space, a combination of both top-down views and ego-centric views, and a human-pushed cart or robot that approximates naturalistic human motion with a socially aware robot. Lastly, we present the TBD Pedestrian Dataset we have collected using our system, which not only surpasses the quantity of similar datasets, but also offers unique pedestrian interaction behavior that adds to the qualitative diversity of pedestrian interaction data.

As mentioned earlier, our approach enables additional data collection in a wide range of locations and constraints. Additional data collection and public updates to this initial dataset are planned. We have also discovered additional challenges with our labeling pipeline on static pedestrians, because static pedestrians have long trajectory duration and constantly adjust their body poses, the resulting trajectories can be noisy and escape the labeler’s attention when using our tool. For future works, we would also explore expanding the diversity of the labels. Some examples include: adding activity labels indicating whether the pedestrian is walking, talking or sitting; adding static obstacle labels for human-object interaction studies; adding group labels for pedestrian groups; and adding gaze direction and head orientation labels for the onboard high-definition 360 camera.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (IIS-1734361), National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR 90DPGE0003), Office of Naval Research (ONR N00014-18-1-2503), and Shimizu Corporation. We would like to thank colleagues in the Tepper School of Business for their assistance on data collection logistics.

REFERENCES

- [1] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, June 2016, pp. 961–971.
- [2] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, June 2018, pp. 2255–2264.
- [3] B. Ivanovic and M. Pavone, "The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs," in *Proc. IEEE/CVF International Conf. on Comput. Vis.*, 2019, pp. 2375–2384.
- [4] K. M. Kitani, B. D. Ziebart, J. A. Bagnell, and M. Hebert, "Activity forecasting," in *Comput. Vis. – ECCV 2012*, 2012, pp. 201–214.
- [5] J. Liang, L. Jiang, K. Murphy, T. Yu, and A. Hauptmann, "The garden of forking paths: Towards multi-future trajectory prediction," in *Proc. IEEE Conf. on Comput. Vis. and Pattern Recognit.*, 2020.
- [6] B. Okal and K. O. Arras, "Learning socially normative robot navigation behaviors with bayesian inverse reinforcement learning," in *2016 IEEE International Conf. on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 2889–2895.
- [7] H. Kretschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.
- [8] A. Biswas, A. Wang, G. Silvera, A. Steinfeld, and H. Admoni, "Socnavbench: A grounded simulation testing framework for evaluating social navigation," *arXiv preprint arXiv:2103.00047*, 2021.
- [9] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone, "Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11807–11814, 2022.
- [10] D. Paez-Granados, Y. He, D. Gonon, D. Jia, B. Leibe, K. Suzuki, and A. Billard, "Pedestrian-robot interactions on autonomous crowd navigation: Reactive control methods and evaluation metrics," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 149–156.
- [11] D. Bršćić, T. Kanda, T. Ikeda, and T. Miyashita, "Person tracking in large public spaces using 3-d range sensors," *IEEE Trans. on Human-Machine Syst.*, vol. 43, no. 6, pp. 522–534, 2013.
- [12] B. Majecka, "Statistical models of pedestrian behaviour in the forum," *Master's thesis, School of Informatics, University of Edinburgh*, 2009.
- [13] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *Comput. Vis. – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., 2016, pp. 549–565.
- [14] S. Oh, A. Hoogs, A. Perera, N. Cuntoor, C.-C. Chen, J. T. Lee, S. Mukherjee, J. Aggarwal, H. Lee, L. Davis, et al., "A large-scale benchmark dataset for event recognition in surveillance video," in *CVPR 2011*. IEEE, 2011, pp. 3153–3160.
- [15] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *Proc. IEEE Int. Conf. Comput. Vis.*, Sept 2009, pp. 261–268.
- [16] A. Lerner, Y. Chrysanthou, and D. Lischinski, "Crowds by example," *Comput. Graph. Forum*, vol. 26, no. 3, pp. 655–664, 2007.
- [17] R. Martin-Martin, M. Patel, H. Rezatofighi, A. Shenoi, J. Gwak, E. Frankel, A. Sadeghian, and S. Savarese, "Jrdb: A dataset and benchmark of egocentric robot visual perception of humans in built environments," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2021.
- [18] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, Oct 2010, pp. 797–803.
- [19] M. Sun, F. Baldini, P. Trautman, and T. Murphey, "Move beyond trajectories: Distribution space coupling for crowd navigation," *arXiv preprint arXiv:2106.13667*, 2021.
- [20] H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, and M. Schwager, "Risk-sensitive sequential action control with multi-modal human trajectory forecasting for safe crowd-robot interaction," in *2020 IEEE/RSJ International Conf. on Intell. Robots and Syst. (IROS)*. IEEE, 2020, pp. 11 205–11 212.
- [21] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core Challenges of Social Robot Navigation: A Survey," *arXiv e-prints*, p. arXiv:2103.05668, Mar. 2021.
- [22] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *CVPR*, 2020.
- [23] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [24] Y. Zhang, P. Sun, Y. Jiang, D. Yu, Z. Yuan, P. Luo, W. Liu, and X. Wang, "Bytetrack: Multi-object tracking by associating every detection box," *arXiv preprint arXiv:2110.06864*, 2021.
- [25] A. Sadeghian, V. Kosaraju, A. Sadeghian, N. Hirose, H. Rezatofighi, and S. Savarese, "Sophie: An attentive gan for predicting paths compliant to social and physical constraints," in *Proc. IEEE/CVF Conf. on Comput. Vis. and Pattern Recognit. (CVPR)*, June 2019.
- [26] A. Mohamed, K. Qian, M. Elhoseiny, and C. Claudel, "Socialstgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction," in *Proc. IEEE/CVF Conf. on Comput. Vis. and Pattern Recognit. (CVPR)*, June 2020.
- [27] A. Wang and A. Steinfeld, "Group split and merge prediction with 3D convolutional networks," *IEEE Trans. Robot. Autom.*, vol. 5, no. 2, pp. 1923–1930, 2020.
- [28] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrilu, and K. O. Arras, "Human motion trajectory prediction: A survey," *arXiv preprint arXiv:1905.06113*, 2019.
- [29] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *IEEE/RSJ International Conf. on Intell. Robots and Syst. (IROS)*, Madrid, Spain, Sept. 2018.
- [30] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *Proc. IEEE International Conf. on Robotics and Automation (ICRA)*, 2019, pp. 6015–6022.
- [31] Y. Chen, C. Liu, B. E. Shi, and M. Liu, "Robot navigation in crowds by graph convolutional networks with attention learned from human gaze," *IEEE Trans. Robot. Autom.*, vol. 5, no. 2, pp. 2754–2761, 2020.
- [32] L. Tai, J. Zhang, M. Liu, and W. Burgard, "Socially compliant navigation through raw depth inputs with generative adversarial imitation learning," in *2018 IEEE International Conf. on Robotics and Automation (ICRA)*, 2018, pp. 1111–1117.
- [33] Y. Gao and C.-M. Huang, "Evaluation of socially-aware robot navigation," *Frontiers in Robotics and AI*, p. 420, 2021.
- [34] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human-robot cooperation," *International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.
- [35] C. Cao, P. Trautman, and S. Iba, "Dynamic channel: A planning framework for crowd navigation," in *2019 International Conf. on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5551–5557.
- [36] A. Wang, C. Mavrogiannis, and A. Steinfeld, "Group-based motion prediction for navigation in crowded environments," in *Conf. on Robot Learning*. PMLR, 2022, pp. 871–882.
- [37] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. Kaesemodel Pontes, D. Ramanan, P. Carr, and J. Hays, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung, Eds., vol. 1. Curran, 2021.
- [38] P. Dendorfer, H. Rezatofighi, A. Milan, J. Shi, D. Cremers, I. Reid, S. Roth, K. Schindler, and L. Leal-Taixé, "Mot20: A benchmark for multi object tracking in crowded scenes," *arXiv:2003.09003[cs]*, Mar. 2020, arXiv: 2003.09003.
- [39] C. Gloor, "Pedsim: Pedestrian crowd simulation," URL <http://pedsim.silmari.org>, vol. 5, no. 1, 2016.
- [40] N. Tsoi, M. Hussein, J. Espinoza, X. Ruiz, and M. Vázquez, "Sean: Social environment for autonomous navigation," in *Proc. 8th International Conf. on Human-Agent Interaction*, 2020, pp. 281–283.
- [41] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core challenges of social robot navigation: A survey," *arXiv preprint arXiv:2103.05668*, 2021.
- [42] M. Kuribayashi, T. Ishihara, D. Sato, J. Vongkulbhisal, K. Ram, S. Kayukawa, H. Takagi, S. Morishima, and C. Asakawa, "Pathfinder: Designing a map-less navigation system for blind people in unfamiliar buildings," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023.

- [43] D. Bršćić, H. Kidokoro, Y. Suehiro, and T. Kanda, "Escaping from children's abuse of social robots," in *Proc. of the tenth annual acm/ieee international Conf. on human-robot interaction*, 2015, pp. 59–66.
- [44] Z. Yan, T. Duckett, and N. Bellotto, "Online learning for human classification in 3d lidar-based tracking," in *2017 IEEE/RSJ International Conf. on Intell. Robots and Syst. (IROS)*. IEEE, 2017, pp. 864–871.
- [45] A. Rudenko, T. P. Kucner, C. S. Swaminathan, R. T. Chadalavada, K. O. Arras, and A. J. Lilienthal, "Thör: Human-robot navigation data collection and accurate motion trajectories dataset," *IEEE Trans. Robot. Autom.*, vol. 5, no. 2, pp. 676–682, 2020.
- [46] T. Chavdarova, P. Baqué, S. Bouquet, A. Maksai, C. Jose, T. Bagautdinov, L. Lettry, P. Fua, L. Van Gool, and F. Fleuret, "Wildtrack: A multi-camera hd dataset for dense unscripted pedestrian detection," in *Proc. IEEE Conf. on Comput. Vis. and Pattern Recognit.*, 2018, pp. 5030–5039.
- [47] B. Benfold and I. Reid, "Stable multi-target tracking in real-time surveillance video," in *CVPR 2011*. IEEE, 2011, pp. 3457–3464.
- [48] B. Zhou, X. Wang, and X. Tang, "Understanding collective crowd behaviors: Learning a mixture model of dynamic pedestrian-agents," in *2012 IEEE Conf. on Comput. Vis. and Pattern Recognit.* IEEE, 2012, pp. 2871–2878.
- [49] A. Alahi, V. Ramanathan, and L. Fei-Fei, "Socially-aware large-scale crowd forecasting," in *Proc. IEEE Conf. on Comput. Vis. and Pattern Recognit.*, 2014, pp. 2203–2210.
- [50] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*. Springer, 2020, pp. 683–700.
- [51] Y. Yuan, X. Weng, Y. Ou, and K. M. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9813–9823.
- [52] K. Mangalam, Y. An, H. Girase, and J. Malik, "From goals, waypoints & paths to long term human trajectory forecasting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 233–15 242.