

Table 2: Map statistics for maps currently supported by trajdata. Note that the INTERACTION dataset does not provide pedestrian walkway or crosswalk information [39].

Dataset	Lane Length (km)	Road Area (m <sup>2</sup> )	Pedestrian Area (m <sup>2</sup> )
nuScenes [18]	212.85	946,275	250,164
INTERACTION [39]	18.78	76,502	–
Lyft Level 5 [19]	185.42	591,333	17,359
nuPlan [27]	325.95	1,327,965	271,277

Table 3: Proportion of stationary agents per dataset. S/M denotes Single/Multi.

Dataset	Proportion	Dataset	Proportion
ETH [22]	4.0%	INTERACTION S/M [39]	5.2%/4.5%
UCY [23]	0.0%	Lyft Level 5 [19]	0.1%
SDD [40]	5.1%	Waymo Open [17]	53.6%
nuScenes [18]	17.5%	nuPlan [27]	0.1%

## 471 A Additional Dataset Details

472 **Licenses.** The ETH [22] and UCY [23] datasets are provided for research purposes<sup>4</sup>. SDD [40] is  
 473 licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 (CC BY-NC-  
 474 SA 3.0) License. nuScenes [18] and nuPlan [27] are mostly licensed under a Creative Commons  
 475 Attribution-NonCommercial-ShareAlike 4.0 International Public License (CC BY-NC-SA 4.0),  
 476 with modifications outlined in <https://www.nuscenes.org/terms-of-use>. The INTERAC-  
 477 TION [39] dataset is provided for non-commercial teaching and research use<sup>5</sup>. The Lyft Level  
 478 5 [19] dataset is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike  
 479 4.0 license (CC-BY-NC-SA-4.0). The Waymo Open Motion Dataset [17] is licensed under its own  
 480 Waymo Dataset License Agreement for Non-Commercial Use<sup>6</sup>.

481 **Protecting Personal Privacy.** All datasets supported by trajdata are captured in public spaces.  
 482 Further, each of the ETH [22], UCY [23], SDD [40], and INTERACTION [39] datasets capture data  
 483 in public spaces from elevated fixed traffic cameras or drones, whose birds-eye viewpoints shield the  
 484 faces of pedestrians and drivers from being collected. The nuScenes [18], Lyft Level 5 [19], Waymo  
 485 Open Motion [17], and nuPlan [27] datasets each preserve privacy by leveraging state-of-the-art  
 486 object detection techniques to detect license plates and faces.

## 487 B Map Statistics

488 As shown in Table 2, current AV dataset maps are very large, spanning multiple neighborhoods in  
 489 different cities. The INTERACTION [39] dataset is a notable exception in magnitude, however, due  
 490 to a drone camera’s limited spatial observation range.

## 491 C Stationary Agents

492 Table 3 summarizes the amount of stationary agents per dataset. As can be seen, nuScenes (17.5%)  
 493 and Waymo Open (53.6%) are comprised of many parked vehicles.

<sup>4</sup>See the statement at the top of <https://icu.ee.ethz.ch/research/datasets.html> and in the “Crowds Data” card of <https://graphics.cs.ucy.ac.cy/portfolio>.

<sup>5</sup>See <http://interaction-dataset.com/terms-for-non-commercial-use> for full terms.

<sup>6</sup>Full terms can be found at <https://waymo.com/open/terms/>.

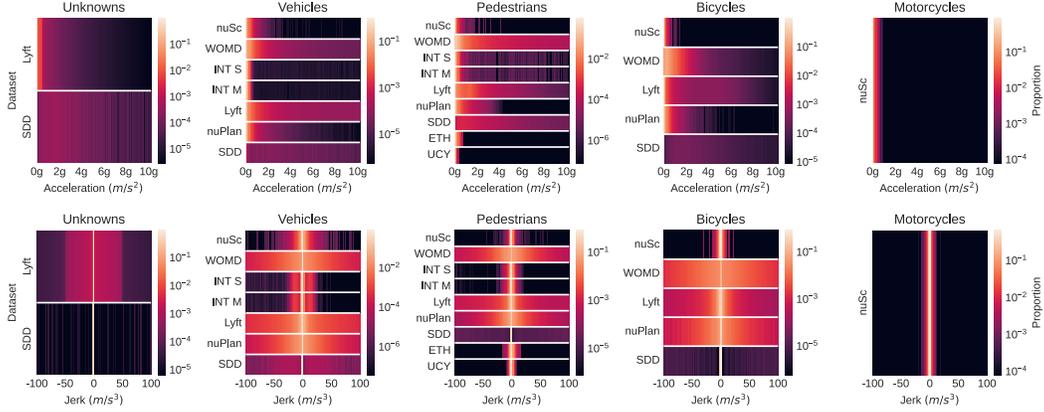


Figure 12: Acceleration (**top**) and jerk (**bottom**) distributions.

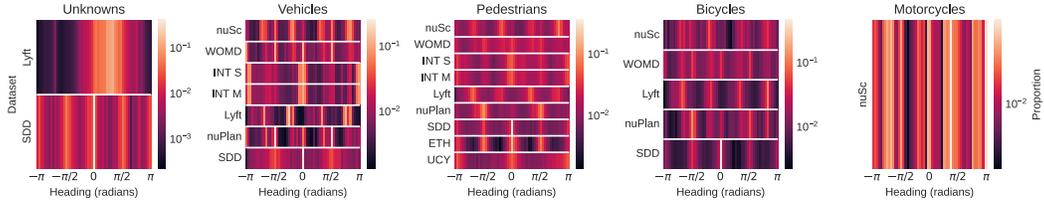


Figure 13: Unnormalized agent heading distributions per dataset and agent type.

## 494 D Acceleration and Jerk

495 Similar to the speed distributions in the main text, Fig. 12 shows that Waymo [17], Lyft [19], and  
 496 SDD [40] have long-tailed acceleration and jerk distributions. Further, as seen in the main text, the  
 497 fixed-camera-based ETH [22] and UCY [23] datasets, as well as the drone-based INTERACTION  
 498 dataset [39], yield tightly-clustered distributions around smooth motion (generally having small  
 499 acceleration and jerk magnitudes). Note that these values are derived by trajdata via finite  
 500 differencing. Accordingly, some overestimation of the acceleration and jerk distribution supports are  
 501 to be expected.

## 502 E Heading Distributions

503 The distributions of (unnormalized) agent headings are shown in Fig. 13. As can be seen, most  
 504 distributions contain peaks around 0 and  $\pm\pi/2$  radians, as north-south and east-west roads are very  
 505 common in many cities. As a particular example, Fig. 14 visualizes heading distributions for vehicles  
 506 and pedestrians in the Waymo Open Motion Dataset [17], showing that pedestrians have much more  
 507 varied heading values than road-following vehicles.

## 508 F Path Efficiency

509 As can be seen in Fig. 15, most path efficiency distributions are uniformly distributed, with peaks  
 510 near 100% (shown as brightly-colored regions), echoing earlier straight-line findings. Further, the  
 511 INTERACTION [39] dataset is an outlier in that its vehicle and pedestrian trajectories are virtually  
 512 all straight lines with much less curved motion than other AV and pedestrian datasets.

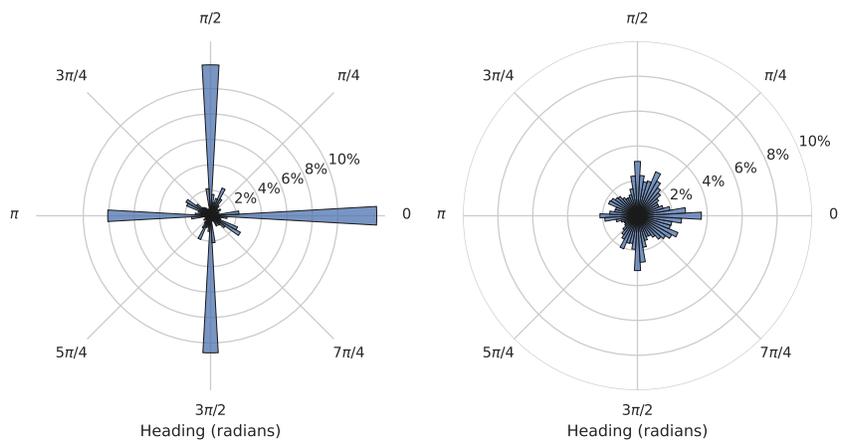


Figure 14: Vehicle (**left**) and pedestrian (**right**) heading distributions in Waymo Open [17].

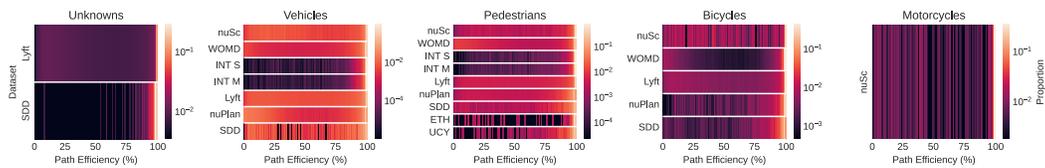


Figure 15: Path efficiency distributions visualized per agent type and dataset.