

OpenTraj: Assessing Prediction Complexity in Human Trajectories Datasets

Javad Amirian¹, Bingqing Zhang², Francisco Valente Castro³, Juan José Baldelomar³, Jean-Bernard Hayet³, and Julien Pettré¹

¹ Univ Rennes, Inria, CNRS, IRISA, France

² University College London, UK

³ CIMAT, A.C., México

Abstract. Human Trajectory Prediction (HTP) has gained much momentum in the last years and many solutions have been proposed to solve it. Proper benchmarking being a key issue for comparing methods, this paper addresses the question of evaluating how complex is a given dataset with respect to the prediction problem. For assessing a dataset complexity, we define a series of indicators around three concepts: Trajectory predictability; Trajectory regularity; Context complexity. We compare the most common datasets used in HTP in the light of these indicators and discuss what this may imply on benchmarking of HTP algorithms. Our source code is released on Github ¹.

Keywords: Human trajectory prediction; trajectory dataset; motion prediction; trajectory forecasting; dataset assessment; benchmarking

1 Introduction

Human trajectory prediction (HTP) is a crucial task for many applications, ranging from self-driving cars to social robots, etc. The communities of computer vision, mobile robotics, and crowd dynamics have been noticeably active on this topic. Many outstanding prediction algorithms have been proposed, from physics-based social force models [1–3] to data-driven models [4–6].

In parallel, efforts have been made towards a proper benchmarking of the existing techniques. This has led to the creation of pedestrians trajectories datasets for this purpose, or to the re-use of datasets initially designed for other purposes, such as benchmarking Multiple Object Tracking algorithms. Most HTP works [3–5, 7] report performance on the sequences of two well-known HTP datasets: the ETH dataset [2] and the UCY dataset [8]. The metrics for comparing prediction performance involve the Average Displacement Error (ADE) and the Final Displacement Error (FDE) on standardized prediction tasks. Other datasets have been used in the same way, but performance comparisons are sometimes subject to controversy, and it remains hard to highlight how significant good performance on a particular sequence or dataset means about a prediction algorithm.

¹ <https://github.com/crowdbot/OpenTraj>




Fig. 1: Taxonomy of trajectories datasets for Human Trajectory Prediction.

In this paper, we address the following questions: (1) How to measure the complexity or difficulty of a particular dataset for the prediction task? (2) How do the currently used HTP datasets compare to each other? Can we draw conclusions about the strengths/weaknesses of state of the art algorithms?

Our contributions are two-fold: (1) We propose a series of meaningful and interpretable indicators to assess the complexity behind an HTP dataset, and (2) we analyze some of the most common datasets through these indicators.

In Section 3, we categorize datasets complexity along three axes, trajectories predictability, trajectories regularity, and context complexity. In Section 4, we define indicators quantifying the complexity factors. In Section 5, we apply these indicators on common HTP datasets and we discuss the results in Section 6.

2 Related work: HTP datasets

Due to the non-rigidness nature of the human body or occlusions, people tracking is a difficult problem and has attracted notable attention. Many video datasets have been designed as benchmarking tools for this purpose and used intensively in HTP. Following the recent progress in autonomous driving, other datasets have emerged, involving more complex scenarios. In this section, we propose a taxonomy of HTP datasets and review some of the most representative ones.

2.1 The zoo of HTP datasets: A brief taxonomy

Many intertwined factors explain how some trajectories or datasets are harder to predict than others for HTP algorithms. In Fig. 1, we summarize essential factors behind prediction complexity, as circles; we separate hidden (blue) and controlled (green) factors. Among hidden factors, we emphasize those related to the acquisition (noisy data), to the environment (multi-modality), or to crowd-related factors (interactions complexity). Some factors can be controlled, such as the recording platform or the choice of the location. To illustrate the variety of setups, snapshots from common HTP datasets are given in Fig. 2.

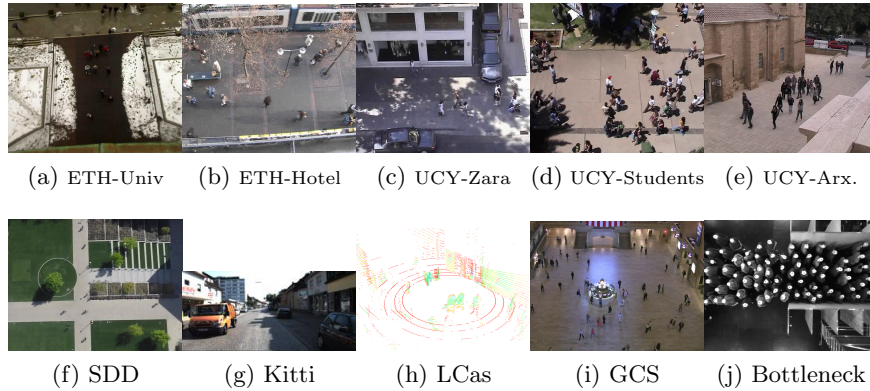


Fig. 2: Sample snapshots from a few common HTP datasets.

Raw data may be recorded by a single [2] or multiple [9] sensors, ranging from monocular cameras [10–12] to stereo-cameras, RGB-D cameras, LiDAR, RADARs, or a mix [13, 9]. Sensors may provide 3D annotations, but most HTP algorithms run on 2D data (the ground plane), and we focus here on 2D analysis.

Annotation is either manual [2, 8, 14], semi-automatic [15], or fully automatic, using detection algorithms [10]. In most datasets, the annotations provide the agents’ positions in the image. Annotated positions can be projected from image coordinates to world coordinates, given homographies or camera projection matrices. For moving sensors (robots [16] or cars [14, 9, 13]), the data are sensor-centered, but odometry data are provided to get all positions in a common frame.

2.2 A short review of common HTP datasets

HTP Datasets from static cameras and drones. The Performance Evaluation of Tracking and Surveillance (PETS) workshops have released several datasets for benchmarking Multiple Object Tracking [17] systems. In particular, the 11 sequences of the PETS’2009 dataset [18], recorded through 8 monocular cameras, include data from *acting* pedestrians, with different levels of density, and have been used in HTP benchmarking [19]. The Town-Centre dataset [12] was also released for visual tracking purposes, with annotations of video footage monitoring a busy town center. It involves around two thousand walking pedestrians with well structured (motion along a street), natural behaviors. The Wild Track dataset [20] was designed for testing person detection in harsh situations (dense crowds) and provides 312 pedestrian trajectories in 400-frame sequences (from 7 views) at 2fps. The EIF dataset [21] gives $\sim 90k$ trajectories of persons in a university courtyard, from an overhead camera. The BIWI pedestrian dataset [2] is composed of 2 scenes with hundreds of trajectories of pedestrians engaged in walking activities. The ATC [22] dataset contains annotations for 92 days of pedestrian trajectories in a shopping mall, acquired from 49 3D sensors.

The UCY dataset [23] provides three scenes with walking/standing activities. Developed for crowd simulation, it exhibits different crowd density levels and a clear flow structure. The Bottleneck dataset [24] also arose from crowd simulation and involved crowd controlled experiments (e.g., through bottlenecks).

VIRAT [25] has been designed for activity recognition. It contains annotated trajectories on 11 distinct scenes, in diverse contexts (parking lot, university campus) and mostly natural behaviors. It generally involves one or two agents and objects. A particular case of activity recognition is the one of sports activities [26], for which many data are available through players tracking technology.

The Stanford Drone Dataset (SDD) [11] is a large scale dataset with 60 sequences in eight scenes, filmed from a still drone. It provides trajectories of $\sim 19k$ moving agents in a university campus, with interactions between pedestrians, cyclists, skateboarders, cars, buses. DUT and CITR [15] datasets have also been acquired from hovering drones for evaluating inter-personal and car-pedestrian interactions. They include, respectively, 1793 and 340 pedestrian trajectories. The inD dataset [10], acquired with a static drone, contains more than 11K trajectories of road users, mostly motorized agents. The scenarios are oriented to urban mobility, with scenes at roundabouts or road intersections. Ko-PER [27] pursues a similar motivation of monitoring spaces shared between cars and non-motorized users. It provides trajectories of pedestrians and vehicles at one road intersection, acquired through laser scans and videos. Similarly, the VRU dataset [28] features around 80 cyclists trajectories, recorded at an urban intersection using cameras and LiDARs. The Forking Paths Dataset [29] was created under the Carla 3D simulator, but it uses real trajectories, which are extrapolated by human annotators to simulate multi-modality with different latent goals.

AV datasets. Some datasets offer data collected for training/benchmarking algorithms for autonomous vehicles (AV). They may be more difficult because of the mobile data acquisition and because the trajectories are often shorter. LCAS [30] was acquired from a LiDAR sensor on a mobile robot. KITTI [14] has been a popular benchmarking source in computer vision and robotics. Its tracking sub-dataset provides 3D annotations (cars/pedestrians) for ~ 20 LiDAR and video sequences in urban contexts. AV companies have recently released their datasets, as Waymo [13], with hours of high-resolution sensor data or Argo AI with its Argoverse [31] dataset, featuring 3D tracking annotations for 11k tracked objects over 113 small sequences. Ntunomy disclosed its nuScenes dataset [9] with 85 annotated scenes in the streets of Miami and Pittsburgh.

Benchmarking through meta-datasets. Meta-datasets have been designed for augmenting the variety of environments and testing the generalization capacities of HTP systems. TrajNet [19] includes ETH, UCY, SDD and PETS; in [32], Becker et al. proposed a comprehensive study over the TrajNet training set, giving tips for designing a good predictor and comparing traditional regression baselines vs. neural-network schemes. Trajnet++ [33] proposes a hierarchy of categorization among trajectories to better understand trajectory distributions within datasets. By mid-2020, over 45 solutions have been submitted on

Trajanet, with advanced prediction techniques [32, 4, 5, 34, 35], but also Social-Force-based models [1], and variants of linear predictors, that give accuracy levels of 94% of the best model [35]. In this work, we give tools to get a deeper understanding of the intrinsic complexities behind these datasets.

3 Problem description and formulation of needs in HTP

3.1 Notations and problem formulation

A trajectory dataset is referred to as \mathbb{X} . We assume that it is made of N_a trajectories of distinct agents. To be as fair as possible in our comparisons, we mainly reason in terms of absolute time-stamps, even though the acquisition frequency may vary. Within \mathbb{X} , the full trajectory of the i -th agent ($i \in [1, N_a]$) is denoted by \mathbf{T}^i , its starting time as τ^i , its duration as δ^i . For $t \in [\tau^i, \tau^i + \delta^i]$, we refer to the state of agent i at t as \mathbf{x}_t^i . We observe \mathbf{x}_t^i only for a finite subset of timestamps (at camera acquisition times). The *frames* are defined as the set of observations at those times and are denoted by \mathbf{F}_t . Each frame contains K_t agents samples.

The state \mathbf{x}_t^i includes the 2D position \mathbf{p}_t^i in a Cartesian system in *meter*. It is often obtained from images and mapped to a world frame; the velocity \mathbf{v}_t^i , in *m/s*, can be estimated by finite differences or filtering.

To compare trajectories, following a common practice in HTP, we split all the original trajectories into N_t trajlets with a common duration $\Delta = 4.8s$. HTP uses trajlets of Δ_{obs} seconds as observations and the next Δ_{pred} seconds as the prediction targets. Hereafter, the set of distinct trajectories of duration Δ obtained this way are referred to as \mathbf{X}^k where $k \in [1, N_t]$ covers the trajlets (with potentially repetitions of the same agent). Typically, $N_t \gg N_a$. Each trajlet may be seen as an observed part and its corresponding target is referred to as \mathbf{X}_+^k .

In the following, we use functions operating at different levels, with different writing conventions. *Trajectory-level functions* $F(\mathbf{X})$, with capital letters, act on trajlets \mathbf{X} . Sometimes, we consider the values of F at specific time values t , at we denote the functions as $F_t(\mathbf{X})$. *Frame-level functions* $\mathcal{F}(\mathbf{F})$ act on frames \mathbf{F} .

3.2 Datasets complexity

We define three families of indicators over trajectory datasets that allow us to compare them and identify what makes them more “difficult” than other.

Predictability. A dataset can be analyzed through how easily individual trajectories can be predicted given the rest of the dataset, independently from the predictor. Low predictability on the trajlet distribution $p(\mathbf{X})$ makes forecasting systems struggle with multi-modal predictive distributions, e.g., at crossroads. In that case, stochastic forecasting methods may be better than deterministic ones, as the latter typically average over the outputs seen in the training data.

Trajectory (ir)regularity. Another dataset characterization is through geometrical and physical properties of the trajectories, to reflect irregularities or deviations to “simple” models. We will use speeds, accelerations for that purpose.

Context complexity. Some indicators evaluate the complexity of the context, i.e., external factors that influence the course of individual trajectories. Typically, crowd density has a strong impact on the difficulty of HTP.

These indicators operate at different levels and may be correlated. For example, complex scenes or high crowdedness levels may lead to geometric irregularities in the trajectories and to lower predictability levels. Finally, even though it is common to *combine* datasets, our analysis is focused on individual datasets.

4 Numerical Assessment of a HTP Dataset complexity

Based on the elements from Section 3, we propose several indicators for assessing a dataset difficulty, most of the kind $F(\mathbf{X}^k)$, defined at the level of trajlets \mathbf{X}^k .

4.1 Overall description of the set of trajlets

To explore the distribution $p(\mathbf{T})$ in a dataset, we first consider the distributions of pedestrian positions at a timestep t . We parametrize each trajlet by fitting a cubic spline $\mathbf{p}_k(t)$ with $t \in [0, 4.8]$. For $t \in [0, 4.8]$, we get 50 time samples $\mathcal{S}(t) = \{\mathbf{p}_k(t), 1 \leq k \leq N_t\}$ and analyze $\mathcal{S}(t)$ through clustering and entropy:

- **Number of Clusters** $M_t(\mathbb{X})$: We fit a Gaussian Mixture Model (GMM) to our sample set using Expectation Maximization and select the number of clusters with the Bayesian Information Criterion [36].
- **Entropy** $H_t(\mathbb{X})$: We get a kernel density estimation of $\mathcal{S}(t)$ (see below in Section 4.2) and use the obtained probabilities to estimate the entropy.

High entropy means that many data points do not occur frequently, while low entropy means that most data points are “predictable”. Similarly, a large number of clusters would require a more complex predictive model. Both indicators give us an understanding of how homogeneous through time are all the trajectories in the dataset.

4.2 Evaluating datasets trajlet-wise predictability

To quantify the trajectory predictability, we use the conditional entropy of the predicted part of the trajectory, given its observed part. Some authors [37] have used alternatively the maximum of the corresponding density. For a trajectory $\mathbf{X}^k \cup \mathbf{X}_+^k$, we define the conditional entropy conditioned to the observed \mathbf{X}^k as

$$H(\mathbf{X}^k) = -E_{\mathbf{X}_+}[\log p(\mathbf{X}_+|\mathbf{X}^k)]. \quad (1)$$

We use kernel density estimation with the whole dataset \mathbb{X} (N_t trajectories) to estimate it. We have N_{obs} observed points during the first Δ_{obs} seconds (trajlet \mathbf{X}_k) and N_{pred} points to predict during the last Δ_{pred} seconds (trajlet \mathbf{X}_k^+). We define a Gaussian kernel K_h over the sum of Euclidean distances between the consecutive points along two trajectories \mathbf{X} and \mathbf{X}' with N points each (in \mathbb{R}^{2N}):

$$K_{h,N}(\mathbf{X}, \mathbf{X}') = \frac{1}{(2\pi h^2)^N} \exp(-\frac{1}{2h^2} \|\mathbf{X} - \mathbf{X}'\|^2), \quad (2)$$

where h is a common bandwidth factor for all the dimensions. We get an approximate conditional density as the ratio of the two kernel density estimates

$$p(\mathbf{X}_+ | \mathbf{X}^k) \approx \frac{\frac{1}{N_t} \sum_{l=1}^{N_t} K_{h, N_{obs} + N_{pred}}(\mathbf{X} \cup \mathbf{X}_+^k, \mathbf{X}^l \cup \mathbf{X}_+^l)}{\frac{1}{N_t} \sum_{l=1}^{N_t} K_{h, N_{obs}}(\mathbf{X}_-^k, \mathbf{X}_-^l)}. \quad (3)$$

Since $K_{h, N_{obs} + N_{pred}}(\mathbf{X}^k \cup \mathbf{X}_+^k, \mathbf{X}^l \cup \mathbf{X}_+^l) = K_{h, N_{obs}}(\mathbf{X}, \mathbf{X}^l) K_{h, N_{pred}}(\mathbf{X}_+, \mathbf{X}_+^l)$, we can express the distribution of Eq. 3 as the following mixture of Gaussian:

$$p(\mathbf{X}_+ | \mathbf{X}^k) \approx \sum_{l=1}^{N_t} \omega_l(\mathbf{X}^k) K_{h, N_{pred}}(\mathbf{X}_+, \mathbf{X}_+^l) \text{ with } \omega_l(\mathbf{X}^k) = \frac{K_{h, N_{obs}}(\mathbf{X}^k, \mathbf{X}^l)}{\sum_{l=1}^{N_t} K_{h, N_{obs}}(\mathbf{X}^k, \mathbf{X}^l)}. \quad (4)$$

For a trajlet \mathbf{X}^k , we estimate $H(\mathbf{X}^k)$ by sampling M samples $\mathbf{X}_+^{(i)}$ from Eq. 4:

$$H(\mathbf{X}^k) \approx -\frac{1}{M} \sum_{m=1}^M \log\left(\sum_{l=1}^{N_t} \omega_l(\mathbf{X}^k) K(\mathbf{X}_+^{(m)}, \mathbf{X}_+^l)\right). \quad (5)$$

4.3 Evaluating trajectories regularity

In this section, we define geometric and statistical indicators evaluating how *regular* individual trajectories \mathbf{X}^k in a dataset may be.

Motion properties. A first series of indicators are obtained through *speed distributions*, where speed is defined as: $s(\mathbf{x}_t) = \|\mathbf{v}_t\|$. At the level of a trajectory \mathbf{X}^k , we evaluate the mean and the largest deviation of speeds along the trajectory

$$S^{avg}(\mathbf{X}^k) = \text{average}_{t \in [\tau^k, \tau^k + \delta^k]} (s(\mathbf{x}_t)) \quad (6)$$

$$S^{rg}(\mathbf{X}^k) = \max_{t \in [\tau^k, \tau^k + \delta^k]} (s(\mathbf{x}_t)) - \min_{t \in [\tau^k, \tau^k + \delta^k]} (s(\mathbf{x}_t)). \quad (7)$$

The higher the speed, the larger the displacements and the more uncertain the target whereabouts. Also, speed variations can reflect on high-level properties such as people activity in the environment or the complexity of this environment.

Regularity is evaluated through accelerations $a(\mathbf{x}_t) \approx \frac{1}{dt} [s(\mathbf{x}_{t+dt}) - s(\mathbf{x}_t)]$. It can reflect the interactions of an agent with its environment according to social-force model [1]: agents typically keep their preferred speed while there is no reason to change it. High accelerations appear when an agent avoids collision or joins a group. We consider the average and maximal accelerations along \mathbf{X}^k

$$A^{avg}(\mathbf{X}^k) = \text{average}_{t \in [\tau^k, \tau^k + \delta^k]} (|a(\mathbf{x}_t)|); \quad A^{max}(\mathbf{X}^k) = \max_{t \in [\tau^k, \tau^k + \delta^k]} (|a(\mathbf{x}_t)|). \quad (8)$$

Non-linearity of trajectories. *Path efficiency* is defined as the ratio of the distance between the trajectory endpoints over the trajectory length:

$$F(\mathbf{X}^k) = \frac{\|p_{\tau^k+\delta^k} - p_{\tau^k}\|}{\int_{t=\tau^k}^{\tau^k+\delta^k} dl}. \quad (9)$$

The higher its value, the closer the path is to a straight line, so we would expect that the prediction task will be “easier” for high values of $F(\mathbf{X}^k)$.

Another indicator is the average angular deviation from a linear motion. To estimate it, we align all trajlets by translating them to the origin of coordinate system and rotating them such that the first velocity is aligned with the x axis:

$$\hat{\mathbf{X}}^k = [\mathbf{R}(-\angle \mathbf{v}_0^k) - \mathbf{p}_0^k] \begin{bmatrix} \mathbf{X}^k \\ 1 \end{bmatrix}^T. \quad (10)$$

Then the deviation of a trajectory \mathbf{X}^k at t and its average value are defined as:

$$D_t(\mathbf{X}^k) = \angle \hat{\mathbf{X}}_t^k \text{ and } D(\mathbf{X}^k) = \text{average}_{t \in [\tau^k, \tau^k+\delta^k]} (D_t(\mathbf{X}^k)). \quad (11)$$

4.4 Evaluating the context complexity

The data acquisition context may impact HTP in different ways. It may ease the prediction by introducing correlations: With groups, it can be easier to predict one’s motion from the other group members. In general, social interactions result into adjustments that may generate non-linearities (and lower predictability).

Collision avoidance is the most basic type of interaction. Higher density resulting into more interactions, this aspect is also evaluated by the density metrics below. However, high-density crowds may ease the prediction (e.g., laminar flow of people). To reflect the intensity of collision avoidance-based interactions, we use the *distance of closest approach* (DCA) [38] at t , for a pair of agents (i, j) :

$$\text{dca}(t, i, j) = \sqrt{\|\mathbf{x}_t^i - \mathbf{x}_t^j\|^2 - (\max(0, \frac{(\mathbf{v}_t^i - \mathbf{v}_t^j)^T (\mathbf{x}_t^i - \mathbf{x}_t^j)}{\|\mathbf{v}_t^i - \mathbf{v}_t^j\|}))^2}, \quad (12)$$

and for a trajlet \mathbf{X}^k (relative to an agent i_k), we consider the overall minimum

$$C(\mathbf{X}^k) = \min_{t \in [\tau^k, \tau^k+\delta^k]} \min_j \text{dca}(t, i_k, j). \quad (13)$$

In [39], the authors suggest that time-to-collision (TTC) is strongly correlated with trajectory adjustments. The TTC for a pair of agents i, j , modeled as disks of radius R , for which a collision will occur when keeping their velocity, is

$$\tau(t, i, j) = \frac{1}{\|\mathbf{v}_t^i - \mathbf{v}_t^j\|^2} [\delta_t^{ij} - \sqrt{(\delta_t^{ij})^2 - \|\mathbf{v}_t^i - \mathbf{v}_t^j\|^2 (\|\mathbf{x}_t^i - \mathbf{x}_t^j\|^2 - 4R^2)}] \quad (14)$$

where $\delta_t^{ij} = (\mathbf{v}_t^i - \mathbf{v}_t^j)^T(\mathbf{x}_t^i - \mathbf{x}_t^j)$. In [39], the authors also proposed quantifying the interaction strength between pedestrians as an energy function of τ :

$$E(\tau) = \frac{k}{\tau^2} e^{-\frac{\tau}{\tau^+}}, \quad (15)$$

with k a scaling factor and τ^+ an upper bound for TTC. Like [39], we estimate the actual TTC probability density between pedestrians (from Eq. 14) over the probability density that would arise without interaction (using the time-scrambling approach of [39]). Then we estimate $E(\tau)$ with Eq. 15. As the range of well-defined values for τ may be small, we group the data into 0.2s intervals and use t-tests to find out the lower bound τ^- when two consecutive bins are significantly different ($p < 0.05$). The upper bound τ^+ is fixed as 3s. TTC and energy interaction are extended for trajlets (only if there exists future collision):

$$T(\mathbf{X}^k) = \min_{t \in [\tau^k, \tau^k + \delta^k]} \min_j \tau(t, i_k, j) \text{ and } E(\mathbf{X}^k) = E(T(\mathbf{X}^k)). \quad (16)$$

Density & distance measures. For a frame \mathbf{F}_t , the *Global Density* is defined as the number of agents per unit area $\mathcal{D}(\mathbf{F}_t) = \frac{K_t}{\mathbf{A}(\mathbb{X})}$, with K_t the number of agents present at t and $\mathbf{A}(\mathbb{X})$ the spatial extent of \mathbb{X} , evaluated from the extreme x, y values. The *Local Density* measures the density in a neighborhood. Plaue et al. [40] infer it with a nearest-neighbour kernel estimator. For a point \mathbf{x}_t ,

$$\rho(\mathbf{x}_t) = \frac{1}{2\pi} \sum_{i=1}^{K_t} \frac{1}{(\lambda d_t^i)^2} \exp\left(-\frac{\|\mathbf{x}_t^i - \mathbf{x}_t\|^2}{2(\lambda d_t^i)^2}\right), \quad (17)$$

with $d_t^i = \min_{j \neq i} \|\mathbf{x}_t^i - \mathbf{x}_t^j\|$ the distance from i to its nearest neighbor and $\lambda > 0$ a smoothing parameter. ρ is used to evaluate a trajlet-wise local density indicator

$$L(\mathbf{X}^k) = \max_{t \in [\tau^k, \tau^k + \delta^k]} \rho(\mathbf{x}_t^{i_k}). \quad (18)$$

5 Experiments

In this section, we analyze some common HTP datasets in the light of the indicators presented in the previous section. In Table 1, we give statistics (location, number of agents, duration...) for the datasets we have chosen to evaluate. We gather the most commonly used in HTP evaluation (ETH, UCY, SDD in particular) and datasets coming from a variety of modalities (static cameras, drones, autonomous vehicles...), to include different species from the zoo of Section 2.1.

For those including very distinct sub-sequences, e.g., ETH, UCY, SDD, inD, and Bottleneck (also denoted by BN in the figures), we split them into their constituting sequences. Also, note that we have focused only on pedestrians (no cyclist nor cars). We also ruled out any dataset containing less than 100 trajectories (e.g., UCY Arxiepiskopi or PETS).

To analyze a dataset \mathbb{X} , we apply systematically the following preprocessing

Table 1: General statistics of assessed datasets. The columns present the type of location where the data is collected, the acquisition means, number of annotated pedestrians, the -rounded- duration (in minute or hour), the total duration of all trajectories, number of trajlets, and percent of non-static trajlets, respectively.

	Dataset	Location	Acquisition	#peds	duration	total dur.	#trajlets	non-static
ETH	Univ	univ entrance	top-view cam	360	13m	1h	823	93%
	Hotel	urban street		390	12m	0.7h	484	66%
UCY	Zara	urban street	top-view cam	489	18m	2.1h	2130	75%
	Students	univ campus		967	11.5m	4.5h	4702	96%
SDD	Coupa	univ campus	drone cam	297	26m	4.5h	5,394	41%
	Bookstore			896	56m	9.5h	11,239	54%
	DeathCircle			917	22.3m	4.2h	8,288	62%
inD	inD-Loc(1)	urban intersection	drone cam	800	180m	7.1h	8302	94%
	inD-Loc(2)			2.1k	240m	18h	21234	95%
Bottleneck	1D Flow _(w=180)	simulated corridor	top-view cam	170	1.3m	1h	940	99%
	2D Flow _(w=160)			309	1.3m	1.5h	1552	100%
	Edinburgh Sep{1,2,4,5,6,10}	univ forum	top-view cam	1.2k	9h	3h	2124	83%
	GC Station	train station	surveillance cam	17k	1.1h	79h	76866	99%
	Wild-Track	univ campus	multi-cam	312	3.3m	1.3h	1215	57%
	KITTI	urban streets	lidar& multi-cam	142	5.8m	0.3h	253	93%
	LCas-Minerva	univ-indoor	lidar	878	11m	4.8h	3553	83%

1. Projection to world coordinates, when necessary.
2. Down-sampling the annotations to a 2-3 fps framerate;
3. Application of a Kalman smoothing with a constant acceleration model;
4. Splitting of the resulting trajectories into trajlets \mathbf{X}^k of length $\Delta = 4.8$ s and filtering out trajlets shorter than 1m.

We finally recall the trajlet-wise indicators we have previously introduced:

Overall description	Entropy $H_t(\mathbb{X}^k)$ and clusters $M_t(\mathbb{X})$ (section 4.1).
Predictability	Cond. entropy $H(\mathbf{X}^k)$ (Eq. 5).
Regularity	Speed $S^{avg}(\mathbf{X}^k)$, $S^{rg}(\mathbf{X}^k)$ (Eq. 7). Acceleration $A^{avg}(\mathbf{X}^k)$, $A^{max}(\mathbf{X}^k)$ (Eq. 8). Efficiency $F(\mathbf{X}^k)$ (Eq. 9). Angular deviation $D(\mathbf{X}^k)$ (Eq. 11).
Context	Closest approach $C(\mathbf{X}^k)$ (Eq. 13). Time-to-collision $T(\mathbf{X}^k)$, energy $E(\mathbf{X}^k)$ (Eq. 16). Local density $L(\mathbf{X}^k)$ (Eq. 18).




Fig. 3: Entropy $H_t(\mathbb{X})$ and number of clusters $M_t(\mathbb{X})$, as described in Section 4.1, at different progression rates t , for a dataset \mathbb{X} . Each dot corresponds to one t .

Overall description of the set of trajlets. For the indicators of Section 4.2, we have chosen $h = 0.5m$ for the Gaussian in the kernel-based density estimation; the number of samples used to evaluate the entropy is $M = 30$; the maximal number of clusters when clustering unconditional or conditional trajectories distributions is 21. In Fig. 3, we plot the distributions of the overall entropy and number of clusters, at different progression rates along the dataset trajectories. Without surprise, higher entropy values are observed for the less structured datasets (without main directed flows) such as SDD or inD. The number of clusters follows a similar trend, indicating possible multi-modality.




Fig. 4: Conditional Entropies $H(\mathbf{X}^k)$.

Predictability indicators. In Fig. 4, we depict the values of $H(\mathbf{X}^k)$, with one dot per trajlet \mathbf{X}^k . Interestingly, excepting the Bottleneck sequences, where high density generates randomness, the support for the entropy distributions are similar among datasets. What probably makes the difference are the tails in these

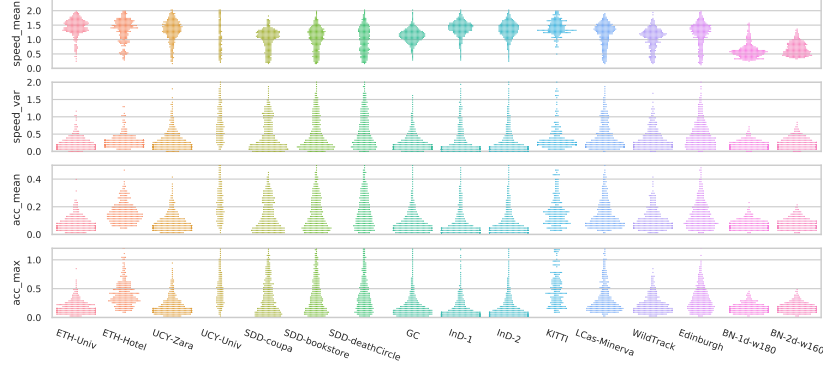



Fig. 5: Speed and acceleration indicators $S^{avg}(\mathbf{X}^k)$, $S^{rg}(\mathbf{X}^k)$, $A^{avg}(\mathbf{X}^k)$, $A^{max}(\mathbf{X}^k)$. From top to bottom: speed means and variations, mean and max. accelerations.


distributions: large lower tails indicate high proportions of easy-to-predict trajlets, while large upper tails indicate high proportions of hard-to-predict trajlets.

Regularity indicators. In Fig. 5, we depict the distributions of the regularity indicators $S^{avg}(\mathbf{X}^k)$, $S^{rg}(\mathbf{X}^k)$, $A^{avg}(\mathbf{X}^k)$, $A^{max}(\mathbf{X}^k)$ from Eqs. 7 and 8. Speed averages are generally centered around 1 and 1.5m/s. Disparities among datasets appear with speed variations and average accelerations: ETH or UCY Zara sequences do not exhibit large speed variations, e.g. compared to Wild Track. In Fig. 6a, we depict the path efficiency $F(\mathbf{X}^k)$ from Eq. 9, and we observe that ETH, UCY paths tend to be straighter. More complex paths appear in Bottleneck, due to the interactions within the crowd, or in SDD-deathCircle, EIF, due to the environment complexity. In Fig. 6b, deviations $D_t(\mathbf{X}^k)$ are displayed for different progression rates along the trajectories, and reflect similar trends.

Context complexity indicators. For estimating the TTC in Eq. 14, we set $R = 0.3m$, and for the interaction energy of Eq. 15, we set $k = 1$. The local density of Eq. 18 uses $\lambda = 1$. In Fig. 7, we display the collision avoidance-related indicators (TTC, DCA and interaction energy) described in Section 4.4, while in Fig. 8, we depict the density-related indicators. Most samples have low interaction energy, but interesting interaction levels are visible Zara, InD. The global density for most datasets stays less than $0.1 p/m^2$ while in InD(1&2), Edinburgh and SDD (Coupa & Bookstore), it is even less than 0.02. Bottleneck (1d & 2d) are significantly high density scenarios. For this reason why we depict them separately. Most natural trajectory datasets have a local density about $0 - 4p/m^2$ while such number is higher ($2 - 4p/m^2$) in Bottleneck. With both density indicators, a dataset such as WildTrack has a high global density and low local density, indicating a relatively sparse occupation. Conversely, low global density and high local density in InD suggests the pedestrians are more clustered. This observation is also reflected in the interaction and entropy indicators as well.



(a) Path Efficiency. The higher, the closer to a straight line.



(b) Deviation from linear motion

Fig. 6: Regularity indicators: Path efficiency and deviation from linear motion.




Fig. 7: Collision avoidance-related indicators: From top to bottom, time-to-collision, Distance of closest approach and Interaction energy.




Fig. 8: Density indicators: On the top, global density (one data point for each frame); on the bottom, local density (one data point for each trajlet).

6 Discussion

Among the findings from the previous Section, Fig. 4 shows that the predictability among most datasets varies in mostly the same ranges. Regarding the motion properties of the datasets (see Fig. 5), another finding is pedestrians’ average speed, which, in most cases, varies from 1.0 to 1.5 m/s. However, this is not the case for *Bottleneck* dataset, because the high density of the crowd does not allow the pedestrians to move with a ‘normal’ speed. In the SDD dataset, we observe multiple pedestrians strolling the campus. As shown in Fig. 6b these low-speed motions are usually associated with high deviation from linear motion, though part of this effect is related to the complexity of the scene layout.

Also, for most of the datasets, the speed variation of trajlets remains almost below 0.5. This is not a true hypothesis for *LCas* and *WildTrack*. As one would expect, the distribution of mean/max acceleration of trajlets is highly correlated with speed variations. In Fig. 6a we see that almost all values are bigger than 90%. For *Bottleneck* we see this phenomenon, where by increasing the crowd density and decreasing crowd speed, the paths become less efficient.

7 Conclusions & Future Work

We have presented in this work a series of indicators for gaining insight into the intrinsic complexity of Human Trajectory Prediction datasets. These indicators cover concepts such as trajectory predictability and regularity, and complexity in the level of inter-pedestrian interactions. In light of these indicators, datasets commonly used in HTP exhibit very different characteristics. In particular, it may explain why predictions techniques that do not use explicit modeling of social interactions, and consider trajectories as independent processes, may be rather successful on datasets where e.g., most trajectories have low collision energy; it may also indicate that some of the more recent datasets with higher levels of density and interaction between agents could provide more reliable information on the quality of the prediction algorithm. Finally, the trajlet-wise analysis presented here opens the door to some evolution in benchmarking processes, as we could evaluate scores by re-weighting the target trajlets in the function of the presented indicators.

Acknowledgements

This research is supported by the CrowdBot H2020 EU Project <http://crowdbot.eu/> and by the Intel Probabilistic Computing initiative. The work done by Francisco Valente Castro was sponsored using an MSc Scholarship given by CONACYT with the following scholar registry number 1000188.

References

1. Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. *Physical review E* **51** (1995) 4282–4286

2. Pellegrini, S., Ess, A., Schindler, K., van Gool, L.: You'll never walk alone: Modeling social behavior for multi-target tracking. In: Proc. of the IEEE Int. Conf. on Computer Vision (ICCV). (2009) 261–268
3. Yamaguchi, K., Berg, A.C., Ortiz, L.E., Berg, T.L.: Who are you with and where are you going? In: Proc. of the IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR). (2011) 1345–1352
4. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social lstm: Human trajectory prediction in crowded spaces. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition (CVPR). (2016) 961–971
5. Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., Alahi, A.: Social gan: Socially acceptable trajectories with generative adversarial networks. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition (CVPR). (2018) 2255–2264
6. Salzmann, T., Ivanovic, B., Chakravarty, P., Pavone, M.: Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. arXiv preprint [abs/2001.03093](#) (2020)
7. Amirian, J., Hayet, J.B., Pettré, J.: Social ways: Learning multi-modal distributions of pedestrian trajectories with GANs. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW). (2019) 2964–2972
8. Lerner, A., Chrysanthou, Y., Lischinski, D.: Crowds by example. Computer Graphics Forum **26** (2007) 655–664
9. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuScenes: A multimodal dataset for autonomous driving. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition (CVPR). (2020) 11621–11631
10. Bock, J., Krajewski, R., Moers, T., Runde, S., Vater, L., Eckstein, L.: The ind dataset: A drone dataset of naturalistic road user trajectories at german intersections. arXiv preprint [abs/1911.07692](#) (2019)
11. Robicquet, A., Sadeghian, A., Alahi, A., Savarese, S.: Learning social etiquette: Human trajectory understanding in crowded scenes. In: Proc. of the European Conf. on Computer Vision, Springer (2016) 549–565
12. Benfold, B., Reid, I.: Guiding visual surveillance by tracking human attention. In: Proc. of the British Machine Vision Conference (BMVC). (2009)
13. Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., Vasudevan, V., Han, W., Ngiam, J., Zhao, H., Timofeev, A., Ettinger, S., Krivokon, M., Gao, A., Joshi, A., Zhao, S., Cheng, S., Zhang, Y., Shlens, J., Chen, Z., Anguelov, D.: Scalability in perception for autonomous driving: Waymo open dataset. arXiv preprint [abs/1912.04838](#) (2019)
14. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: the KITTI dataset. The International Journal of Robotics Research **32** (2013) 1231–1237
15. Yang, D., Li, L., Redmill, K., Ozguner, U.: Top-view trajectories: A pedestrian dataset of vehicle-crowd interaction from controlled experiments and crowded campus. In: Proc. of the IEEE Intelligent Vehicles Symposium (IV). (2019) 899–904
16. Ess, A., Leibe, B., Van Gool, L.: Depth and appearance for mobile scene analysis. In: Proc. of the IEEE Int. Conf. on Computer Vision (ICCV). (2007) 1–8
17. Leal-Taixé, L., Milan, A., Reid, I., Roth, S., Schindler, K.: Motchallenge 2015: Towards a benchmark for multi-target tracking. arXiv preprint [abs/1504.01942](#) (2015)
18. Ferryman, J., Shahrokni, A.: Pets2009: Dataset and challenge. In: Proc. of the IEEE Int. Workshop on Performance Evaluation of Tracking and Surveillance (PETS). (2009) 1–6

19. Sadeghian, A., Kosaraju, V., Gupta, A., Savarese, S., Alahi, A.: Trajnet: Towards a benchmark for human trajectory prediction. arXiv preprint **abs/1805.07663** (2018)
20. Chavdarova, T., Baqu, P., Bouquet, S., Maksai, A., Jose, C., Bagautdinov, T., Lettry, L., Fua, P., Van Gool, L., Fleuret, F.: Wildtrack: A multi-camera hd dataset for dense unscripted pedestrian detection. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition (CVPR). (2018) 5030–5039
21. Majecka, B.: Statistical models of pedestrian behaviour in the forum. Master’s thesis, School of Informatics, University of Edinburgh (2009)
22. Brscic, D., Kanda, T., Ikeda, T., Miyashita, T.: Person position and body direction tracking in large public spaces using 3d range sensors. IEEE Transactions on Human-Machine Systems **43** (2013) 522–534
23. Lerner, A., Chrysanthou, Y., Lischinski, D.: Crowds by example. Computer Graphics Forum **26** (2007) 655–664
24. Seyfried, A., Passon, O., Steffen, B., Boltes, M., Rupprecht, T., Klingsch, W.: New insights into pedestrian flow through bottlenecks. Transportation Science **43** (2009) 395–406
25. Oh, S., Hoogs, A., Perera, A., Cuntoor, N., Chen, C.C., Lee, J.T., Mukherjee, S., Aggarwal, J., Lee, H., Davis, L., Swears, E., Wang, X., Ji, Q., Reddy, K., Shah, M., Vondrick, C., Pirsiavash, H., Ramanan, D., Yuen, J., Torralba, A., Song, B., Fong, A., Roy-Chowdhury, A., Desai, M.: A large-scale benchmark dataset for event recognition in surveillance video. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition (CVPR). (2011) 3153–3160
26. Harmon, M., Lucey, P., Klabjan, D.: Predicting shot making in basketball learnt from adversarial multiagent trajectories. arXiv preprint **abs/1609.04849** (2016)
27. Strigel, E., Meissner, D., Seeliger, F., Wilking, B., Dietmayer, K.: The ko-per intersection laserscanner and video dataset. In: Proc. of the IEEE Conf. on Intelligent Transportation Systems (ITSC). (2014) 1900–1901
28. Bieshaar, M., Zernetsch, S., Hubert, A., Sick, B., Doll, K.: Cooperative starting movement detection of cyclists using convolutional neural networks and a boosted stacking ensemble. arXiv preprint **abs/1803.03487** (2018)
29. Liang, J., Jiang, L., Murphy, K., Yu, T., Hauptmann, A.: The garden of forking paths: Towards multi-future trajectory prediction. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition. (2020) 10508–10518
30. Yan, Z., Duckett, T., Bellotto, N.: Online learning for human classification in 3d lidar-based tracking. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS). (2017) 864–871
31. Chang, M.F., Lambert, J.W., Sangkloy, P., Singh, J., Bak, S., Hartnett, A., Wang, D., Carr, P., Lucey, S., Ramanan, D., Hays, J.: Argoverse: 3d tracking and forecasting with rich maps. In: Proc. of the Int. Conf. on Computer Vision and Pattern Recognition (CVPR). (2019) 8748–8757
32. Becker, S., Hug, R., Hbner, W., Arens, M.: An evaluation of trajectory prediction approaches and notes on the trajnet benchmark. arXiv preprint **abs/1805.07663** (2018)
33. Kothari, P., Alahi, A.: Trajnet++: Large-scale trajectory forecasting benchmark. In: Proc. of the 2nd Workshop on Long-Term Human Motion Prediction (ICRA). (2020)
34. Ellis, D., Sommerlade, E., Reid, I.: Modelling pedestrian trajectory patterns with gaussian processes. In: Proc. of the IEEE Int. Conf. on Computer Vision Workshops (ICCVW). (2009) 1229–1234

35. Giuliari, F., Hasan, I., Cristani, M., Galasso, F.: Transformer networks for trajectory forecasting. arXiv preprint **abs/2003.08111** (2020)
36. Claeskens, G., Hjort, N.L. Cambridge Series in Statistical and Probabilistic Mathematics. In: The Bayesian information criterion. Cambridge University Press (2008) 7098
37. Li, M., Westerholt, R., Fan, H., Zipf, A.: Assessing spatiotemporal predictability of lbsn: A case study of three foursquare datasets. *GeoInformatica* **22** (2016)
38. Olivier, A.H., Marin, A., Crtual, A., Berthoz, A., Pettr, J.: Collision avoidance between two walkers: Role-dependent strategies. *Gait and Posture* **38** (2013) 751 – 756
39. Karamouzas, I., Skinner, B., Guy, S.J.: Universal power law governing pedestrian interactions. *Phys. Rev. Lett.* **113** (2014) 238701
40. Plaue, M., Chen, M., Brwolff, G., Schwandt, H.: Trajectory extraction and density analysis of intersecting pedestrian flows from video recordings. In: Proc. of the ISPRS Conf. on Photogrammetric Image Analysis. (2011) 285–296