

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LEARNING TO DEAGGREGATE: LARGE-SCALE TRAJECTORY GENERATION WITH SPATIAL PRIORS

006 **Anonymous authors**

007 Paper under double-blind review

010 ABSTRACT

013 Generating realistic large-scale trajectories is essential for applications in urban
014 mobility and transportation, yet current generative models either do not offer any
015 controllability or rely on strong sample-specific conditioning. We introduce the
016 Temporal Deaggregation Diffusion Model (TDDM), a hierarchical framework that
017 first represents mobility using spatial priors, which are marginal distributions over
018 geographical occupancy, and then deaggregates them into trajectories. This separa-
019 tion enables generation without sample-specific conditions, supporting transfer
020 to new regions. To support evaluation, we build a benchmark across three cities
021 spanning different continents (Asia, Europe, North America), with standardized
022 metrics for fidelity and distributional coverage. Across all datasets, TDDM im-
023 proves trajectory fidelity and coverage over leading baselines, and demonstrates
024 stable performance when applied to unseen cities. By explicitly decoupling spa-
025 tial allocation from temporal realization, our work highlights the role of spatial
026 occupancy priors in enabling scalable and generalizable trajectory generation.

027 1 INTRODUCTION

029 Time-series data of human mobility enables applications such as pandemic forecasting and manage-
030 ment (Ilin et al., 2021), smart city development (Wang et al., 2022), urban governance (Xiong et al.,
031 2024), human rights violation detection (Tai et al., 2022) and monitoring of global migration induced
032 by war and climate change (Niva et al., 2023; Alessandrini et al., 2020). Two major challenges stand
033 in the way for using time-series data for these purposes.

034 The first is a *shortage of publicly available data* (Ansari et al., 2024). Data can only be collected and
035 shared in limited capacity due to concerns of privacy, business and national security, creating a silo
036 effect. Secondly, *generalization beyond observed data*, such as to new regions, unseen spatial areas,
037 or rapidly changing environments, is often a necessary complement to the readily collectible data.
038 One such case is generating high-fidelity realistic spatio-temporal trajectory data, such as individual
039 pedestrians navigating a city or a building. Open problems within the road traffic domain (Lana et al.,
040 2018), and using human mobility data at large, are (1) high quality large-scale trajectories and (2)
041 adaptation to sudden environmental changes. Both are hindered by the unavailability of data, either
042 because existing data cannot be shared or because new environments lack sufficient observations.

043 A promising direction is to use time-series generative models to capture and generalize mobility
044 distributions. Although these models can be adapted for privacy (Yoon et al., 2019b; Wang et al., 2023;
045 Buchholz et al., 2024) or forecasting (Alcaraz & Strothoff, 2023), this work focuses exclusively on
046 improving fidelity and cross-region generalization.

047 Existing approaches, while promising, either fail to capture the multi-modal structure of mobility
048 data or struggle to scale across diverse environments (Buchholz et al., 2024). GAN- and VAE-
049 based approaches such as TimeGAN (Yoon et al., 2019a), TimeVAE (Desai et al., 2021), COSCI-
050 GAN (Seyfi et al., 2022), and TrajGen (Cao & Li, 2021) suffer from mode collapse and oversimplified
051 representations, with unconditional methods additionally offering no control over generated patterns.
052 Recent diffusion models improve fidelity but either remain unconditional (Diffusion-TS (Yuan & Qiao,
053 2024)) or rely on strong sample-specific conditioning (DiffTraj (Zhu et al., 2023), ControlTraj (Zhu
et al., 2024)), limiting generalization across regions. While sample-specific conditioning increases

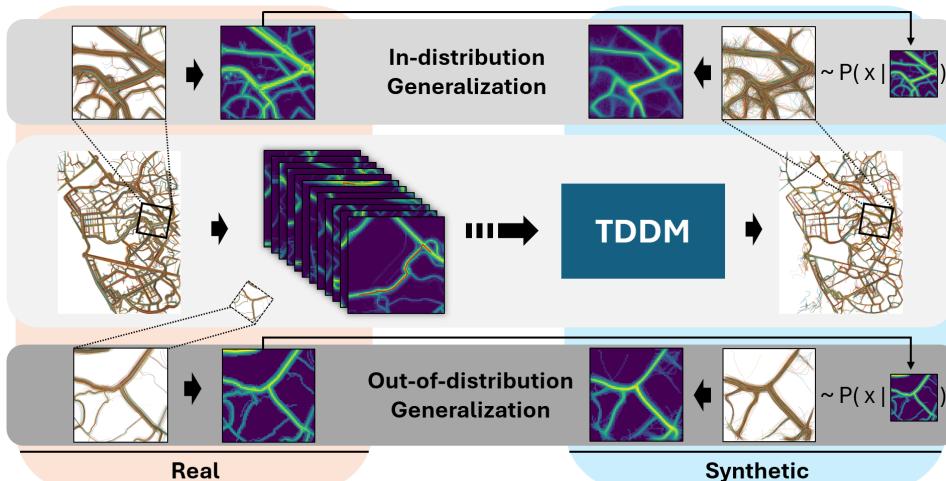


Figure 1: TDDM is trained on real 2D trajectories (left) to generate synthetic trajectories (right), conditioned on how likely it should be for the population of synthetic trajectories to occupy space on the 2D plane (spatial prior). The latter is represented as a discrete distribution over occupancy frequency, i.e., the marginal distribution over the trajectory probability distribution if integrating out time. This yields both high-fidelity in-distribution generalization (top) and out-of-distribution generalization (bottom), the latter when conditioning on a marginal distribution not part of the training data (dashed rectangle). Trajectory data from Beijing, China (Geolife).

quality of the synthetic trajectories, it increases the risk of memorization and prevents cross-region generalization by potentially tying each sample to a training example. Our key insight is that trajectory generation can be factorized into two components: *where* people move, encoded as spatial priors, and *how* people move temporally.

To this end, we introduce Temporal Deaggregation Diffusion Model (TDDM), which conditions on spatial aggregate priors rather than trajectory-level statistics, combining controllability with high fidelity and cross-region generalization. We capture spatial patterns through marginal distributions over local regions, then condition a diffusion model to generate temporal trajectories that respect these patterns. Critically, we canonicalize each region before modeling via a similarity transform, enabling a single model to generalize across all regions for a given scale (e.g. 3x3 km). This spatial-temporal separation enables city-to-city generalization: the model learns temporal dynamics that are invariant to absolute location and orientation. Figure 1 illustrates capabilities of TDDM in a setting of mobility trajectory data. Figure 2 provides a visual comparison with baseline methods.

Our main contributions are:

- *Spatial-Temporal Factorization*: We propose TDDM, a diffusion-based trajectory model that factorizes generation into spatial occupancy priors and temporal dynamics, with coordinate normalization enabling parameter sharing across geographic regions.
- *Benchmarking at Scale*: We establish a standardized evaluation framework across three cities on different continents (Beijing, Porto, San Francisco), with trajectory-specific metrics that harmonize sample fidelity, distributional coverage, and downstream usefulness.
- *Improved Fidelity and Coverage*: TDDM consistently outperforms leading baselines on KL-based distributional measures, demonstrating improved support coverage and proportionality while maintaining strong fidelity across datasets.
- *Generalization to New Regions*: Leveraging spatial priors and canonicalization, TDDM generates realistic trajectories in unseen parts of a city and in entirely new cities without retraining or finetuning, showing strong out-of-distribution zero-shot performance.

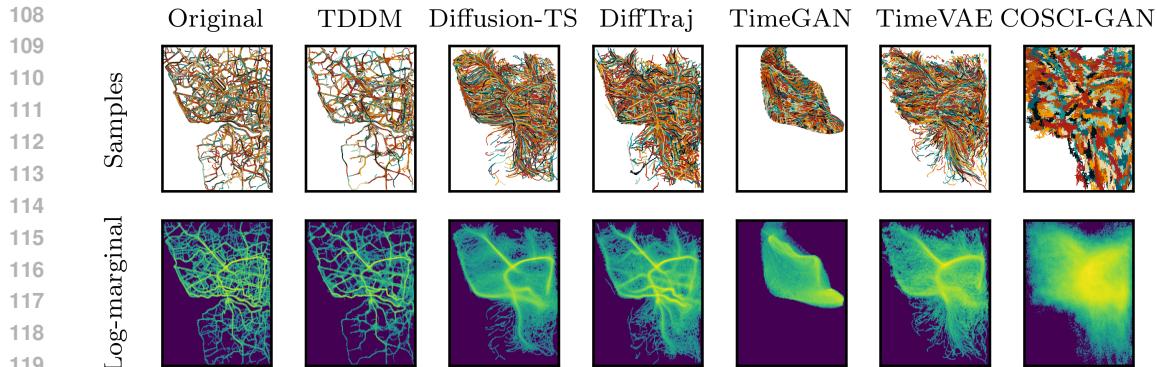


Figure 2: Comparison between original and synthetic trajectories for Porto dataset. First row shows individual trajectory samples while bottom shows log-density heatmaps of all observations. The synthetic data of the proposed model, TDDM, (second column) most closely matches both the individual trajectory patterns and overall density distribution of the original data.

2 PROBLEM DEFINITION

The task of learning an unconditional generative model is defined as learning a mapping f from samples drawn from a known distribution $\mathcal{D}_{\text{known}}$, e.g. standard normal distribution, to samples from an unknown target distribution $\mathcal{D}_{\text{unknown}}$. The mapping is learned without direct access to the unknown distribution, and is instead limited to a set of samples $\mathbb{X}_{\text{train}} = \{x_1, \dots, x_N\}$, where $x_i \sim \mathcal{D}_{\text{unknown}}$. Once the mapping has been learned, synthetic data can be generated by first sampling the known distribution, then passing the individual samples through the mapping function $\mathbb{X}_{\text{synthetic}} = \{f(y_i)\}_{i=1}^M$, where $y_i \sim \mathcal{D}_{\text{known}}$. The goal of this mapping is for the synthetic samples to be *similar* to samples from the known distribution, both on a sample level and as a distribution.

3 TEMPORAL DEAGGREGATION DIFFUSION MODEL

In this work, we focus on the issue of learning and controlling distribution properties of the synthetic data distribution. More specifically, we focus on controlling the spatial marginal of trajectories, i.e. where time has been marginalized out. This is achieved by conditioning the generative process on a description of the marginal distribution, a spatial prior, with the goal being to generate samples that in aggregate follows said marginal distribution.

The central idea of TDDM is to separate *where* people move from *how* they move in time. We achieve this by conditioning trajectory generation on spatial priors H , which describe marginal occupancy distributions over regions. By factoring spatial allocation from temporal dynamics, TDDM learns invariant motion patterns that generalize across cities.

Concretely, the approach proceeds in three steps: (i) partition the spatial domain into regions, (ii) canonicalize each region via a similarity transform, and (iii) compute region-specific spatial priors that describe the local marginal distribution.

Partitioning the spatial domain. We adopt a hierarchical factorization by partitioning the spatial domain into discrete regions, each associated with a spatial prior. The prior describes the marginal distribution of occupancy in that region after marginalizing out time. This setup enables generation at the region level, where local priors act as controllable conditions for the diffusion process.

For training, the partitioning is into regions of the same shape but with randomized translation and rotation and consequently have arbitrary overlap. For sampling of the trained model, the partitioning can be on a grid which covers the area of interest (the trained-on city, the trained-on part of a city concatenated with the rest of the city, or a new city entirely) with partial border overlap.

162 **Canonicalization.** Each region is mapped into a canonical frame of reference using a similarity
 163 transform (Goodall, 1991), paralleling Procrustes alignment in shape analysis. Unlike group-
 164 equivariant architectures (Cohen & Welling, 2016), which encode invariances into model structure,
 165 our approach achieves invariance via input-output transformation, keeping the architecture lightweight
 166 and without additional inductive bias.

167 For a region r_c , we apply the similarity transformation $T_{r_c}(p, \alpha, s)$, parametrized by translation
 168 $p = -\text{pos}(r_c)$, rotation $\alpha = -\text{rot}(r_c)$, and scaling $s = 2/\text{width}(r_c)$. This maps trajectories and
 169 priors into a normalized coordinate system $[-1, 1]^D$.

170 For example, consider a 1×1 km region in downtown Beijing: T_{r_c} translates the region to the origin,
 171 rotates it to a fixed orientation, and rescales coordinates to $[-1, 1] \times [-1, 1]$. This normalization
 172 enables the model to learn local trajectory patterns (e.g., vehicles turning at intersections) that transfer
 173 across locations and cities. When sampling a region, the inverse transform is applied to each sample.

174 **Spatial priors.** More formally, let $x \in \mathbb{X}$ be a sample consisting of several observations $x[n] \in$
 175 $\mathbb{R}^D, \forall n$, where $D = 2$ denotes spatial dimensions (long, lat). We generate trajectories region by
 176 region. For each subregion r_c , we compute its spatial prior H and express the generative model as:

$$179 \quad p(x) = \int p(x|H)p(H) dl. \quad (1)$$

180 In practice, we set $p(H) = p(H = f(r_c, \mathbb{X})) = p(r_c)$ where H is a discrete marginal distribution,
 181 $\sum_i H_{i,j} = 1$, and r_c is a subregion of r . The probability for a subregion r_c is:

$$184 \quad p(r_c) \propto \sum_{x \in \mathbb{X}} \sum_n \mathbb{1}(T_{r_c}x[n] \in [-1, 1]^D). \quad (2)$$

186 Within each region, the prior is discretized by cells (i, j) :

$$188 \quad H_{i,j} = f(r_c, \mathbb{X})_{i,j} = \frac{\sum_{x \in \mathbb{X}} \sum_n \mathbb{1}_{r_c i,j}(x[n])}{\sum_{x \in \mathbb{X}} \sum_n \mathbb{1}(x[n] \in \mathcal{R}_{r_c})}, \text{ where} \quad (3)$$

$$190 \quad \mathbb{1}_{r_c i,j}(x[n]) = \begin{cases} 1, & \text{if } x[n] \text{ falls within cell } (i, j) \text{ of region } r_c \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

192 In practice, we use a 64×64 grid for 3×3 km regions, balancing spatial detail with computational
 193 efficiency: finer grids increase quadratic token cost in the transformer, while coarser resolution would
 194 reduce the spatial information needed to capture detailed road structure and trajectory patterns. The
 195 final approximation for $p(x)$ is then

$$197 \quad p(x) = \sum_{r_c} p(x|H = f(r_c, \mathbb{X}))p(r_c). \quad (5)$$

199 which is a generative mixture model over region partitions.

200 The spatial prior H provides the context that separates *where* people move from *how* they move. By
 201 marginalizing out time, H encodes only the spatial occupancy of a region. The diffusion model then
 202 learns to generate realistic temporal trajectories that, in aggregate, match this spatial pattern. Because
 203 H can be estimated (even in unseen cities) for new regions (even in unseen cities), while temporal
 204 dynamics remain transferable, this factorization supports cross-region generalization. See Figure 1
 205 for examples of spatial priors H and corresponding trajectories.

206 To learn $p(x|H)$ we propose an architecture based on the denoising diffusion architecture (Ho et al.,
 207 2020), using a transformer encoder (Vaswani et al., 2017) for denoising.

209 Generating synthetic data using denoising diffusion is achieved by learning to reverse a noise-adding
 210 process. See Appendix C.1 for more details on denoising diffusion. There are several ways to
 211 parameterize the denoising process, we extend the noise prediction parameterization (Ho et al., 2020)
 212 to include the marginal distribution: $\epsilon_\theta(x_t, t, H)$. This means that at any step of the denoising process,
 213 we have full access to a discretized version of the marginal distribution of the distribution we are
 214 sampling, a noisy trajectory, as well as the expected noise level via the denoising step.

215 This also poses a challenge, as we need a model that can handle these different modalities. To this
 216 end, we employ a transformer encoder. This allows us to use different strategies to tokenize the

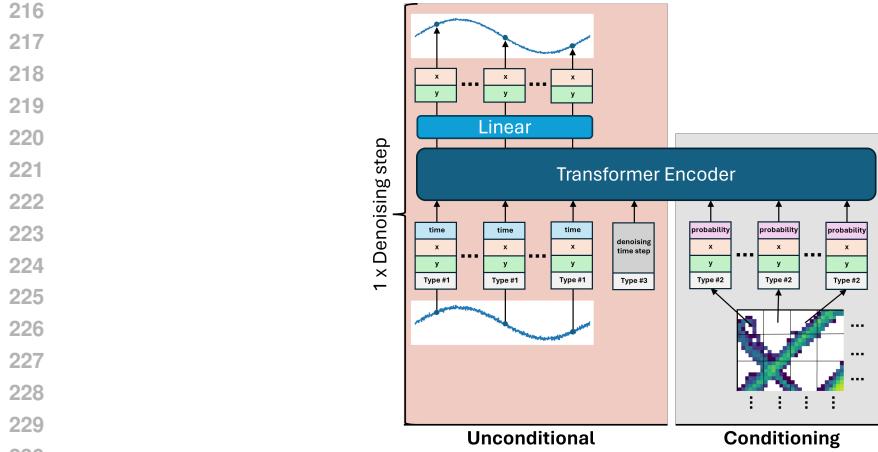


Figure 3: Architecture overview. In the unconditional part, each time point of the noisy trajectory is converted into a separate token with positional encoding used to encode its values and the time point, and a learned vector representing its type. The denoising step token, includes the denoising step encoded with positional encoding and then concatenated with a learned type vector. The input also include a marginal distribution to guide the denoising process to generate samples with particular properties, improving in-distribution performance as well as enabling generalization to previously unobserved areas. The marginal distribution is split into tokens, concatenated with a learned type vector and its position using positional encoding.

different modalities, the resulting tokens are then given as an input to the transformer. The attention mechanism allows the model to learn relevant relationship between the different token types, across different physical distances, as well as in time, to successfully denoise a trajectory.

More specifically, we employ three types of tokens: (1) *trajectory tokens*, (2) *marginal distribution tokens* and (3) *denoising noise step tokens*. At the end of each token, there is a learned token type which allows the model to distinguish between the different types. Trajectories x_i are split into L tokens, one for each time step, and positional encoding is used to encode position and time. For the discretized marginal distribution H , we follow Vision Transformers (Dosovitskiy et al., 2020) and split marginal distribution into several regions, linearly project, and finally prepend the position of the region. Finally, a single token for t using positional encoding is used to denote the current step in the denoising process.

Rather than using a class-token, we use the output corresponding to the L trajectory tokens to make the noise prediction. These are linearly transformed, each token down to 2 dimensions: one for x and one for y , and then concatenated across time to construct the final prediction for the amount of noise added at diffusion step t .

An overview of the architecture is shown in Figure 3, which illustrates how a single trajectory is tokenized, how the marginal distribution is split into subregions and which tokens are used as output. The training and sampling procedures are formalized in Algorithms 1 and 2. Algorithm 1 shows how we train the model by randomly sampling regions, computing their spatial priors, and learning to generate trajectories matching those priors. Algorithm 2 shows how we generate city-scale datasets by conditioning on spatial priors alone to enable zero-shot transfer to new regions or cities. The training and sampling procedures are shown in Algorithm 1 and 2, respectively. See Appendix C for hyperparameters, as well as additional details on tokenization.

To understand how these algorithms enable zero-shot generalization in our out-of-distribution experiments, consider the key design choices. In Algorithm 1, lines 2-3 randomly sample regions and compute spatial priors, ensuring the model encounters diverse geographic contexts during training. Line 6 canonicalizes trajectories to normalized coordinates $[0, 1]^D$, allowing the model to learn location-invariant dynamics. The denoising process (line 10) is conditioned on H , teaching the model to respect spatial distributional constraints. In Algorithm 2, zero-shot transfer works as follows: for target regions (unseen city areas or entirely different cities), we compute the spatial prior H from

270 $\mathbb{X}_{\text{target}}$ trajectories (line 3), but the model ϵ_θ never receives individual target trajectories, only their
 271 aggregate spatial distribution. The model generates trajectories in normalized space conditioned
 272 solely on H (lines 7-10), then transforms them back to global coordinates of the target region (line
 273 11). This canonicalization symmetry (normalize during training, denormalize during generation)
 274 combined with conditioning on aggregate distributions rather than trajectory instances enables the
 275 model to apply learned dynamics to any geographic region without gradient updates or fine-tuning on
 276 target data, demonstrating zero-shot generalization.

277

278

Algorithm 1 Training

280 **Input:** Dataset $\mathbb{X}_{\text{train}}$, Region distribution $p(r_c)$, Target
 281 get length
 282 **Output:** Trained model parameters θ

283 1: **repeat**
 284 2: $r_c \sim p(r_c)$
 285 3: $H \leftarrow f(r_c, \mathbb{X}_{\text{train}})$
 286 4: $\mathcal{X}_{r_c} \leftarrow$ Find contiguous subsequences of trajectories in \mathbb{X} that lie within r_c
 287 5: Filter \mathcal{X}_{r_c} by minimum length, maximum time
 288 gaps, and speed limits
 289 6: Normalize \mathcal{X}_{r_c} to $[0, 1]^D$ relative to r_c
 290 7: Select random sequence $x \in \mathcal{X}_{r_c}$
 291 8: $t \sim \text{Uniform}(\{0, \dots, T\})$
 292 9: $\epsilon \sim N(0, \mathbf{I})$
 293 10: Update model parameters according to:

$$\nabla_\theta ||\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, H, t)||$$

 294 11: **until** converged

295

296

297

298

299

300

4 EVALUATION

301

302 To evaluate TDDM’s trajectory generation capabilities, we address three key questions:

303

1. Does TDDM generate higher-quality synthetic trajectories than existing approaches?
2. Can the spatial-temporal factorization enable better coverage of complex urban distributions?
3. Does coordinate normalization actually enable generalization to unseen regions?

307

308 We start with a comparative study on the unconditional generation task and compare against leading
 309 methods representing major generative paradigms: TimeGAN (seminal time-series GAN), DiffTraj
 310 (UNet-based trajectory diffusion model), Diffusion-TS (transformer-based time-series diffusion
 311 model), TimeVAE (VAE for sequential data), and COSCI-GAN (multi-channel time-series model).
 312 These baselines provide comprehensive coverage of trajectory generation approaches. See Ap-
 313 pendix A for detailed comparisons. There we also include theoretical comparisons with recent work
 314 without reproducible source code (TrajGen Cao & Li (2021), ControlTraj Zhu et al. (2024)) as well
 315 as with COLA Wang et al. (2024).

316

317 Then, we investigate how well TDDM can generalize to new unseen environments. First by training
 318 on a limited part of the map and, using only spatial prior, generating for the remaining map. Second
 319 by showing how a TDDM trained on one city from one dataset can generalize to another city in
 another dataset.

320

321 The evaluation is across three large-scale human trajectory datasets and across several evaluation
 322 measures, which together span current challenges and wanted properties of trajectory generation.

323

Evaluation Measures Several perspectives exists on synthetic data quality in the literature (Alaa
 et al., 2022; Wu et al., 2021; Esteban et al., 2017) covering a range of partially-overlapping aspects.

In this work, we propose to harmonize these and consequently span high fidelity, covering the support, moving beyond copying the training data, downstream task informativeness and distributional proportionality, as captured by the following five qualities: (I) *Fidelity*: (Alaa et al., 2022) The individual synthetic samples should have similar characteristics to, or be indistinguishable from, samples from the original distribution. (II) *Diversity*: (Alaa et al., 2022) It should be possible for synthetic data to be drawn from any part of the unknown distribution’s support. (III) *Proportionality*: (Wu et al., 2021) The probability of a sample occurring in the synthetic distribution should be proportional to the probability of a sample occurring in the unknown distribution. (IV) *Usefulness*: (Esteban et al., 2017) The synthetic data should capture aspects of the unknown distribution that is useful for downstream tasks. (V) *Generalization*: (Alaa et al., 2022) Synthetic samples should not be mere copies of the training data.

As there is not a single measure that can encapsulate all qualities, we propose to use the following evaluation measures to paint a fuller picture of the synthetic data quality:

- **TSTR:** Train on synthetic, test on real (Esteban et al., 2017). A transformer decoder is trained on the synthetic data, with the task to predict 10 future states given a sequence of observations. The resulting model is then evaluated on the training data and the mean absolute value is reported. This evaluates the usefulness, fidelity and diversity of the synthetic data. Lower is better.
- **KL Divergence:** Evaluates support coverage and proportionality between distributions. $KL(R\|S)$ measures how well the synthetic distribution covers areas where real data exists, penalizing missed modes in the real distribution. $KL(S\|R)$ measures how well the synthetic distribution stays within regions supported by real data, penalizing unrealistic generations. Together, they provide complementary perspectives: low $KL(R\|S)$ indicates good coverage of the real distribution, while low $KL(S\|R)$ indicates high fidelity of the synthetic samples. Symmetric KL and Jensen–Shannon divergence balance these concerns, with JS providing better stability when handling regions with zero probability. Lower is better.

We also make use of the measures **Density Error**, **Trip Error**, **Length Error** and **Pattern Score**, introduced by Zhu et al. (2023). See Appendix E for details on all six measures.

Dataset and Preprocessing For evaluation we use three publicly accessible real-world GPS-trajectory datasets: Geolife (Zheng et al., 2011), Porto (Moreira-Matias et al., 2013) and Cabspotting (Piorkowski et al., 2022). Geolife consists of data from 182 users performing various activities (e.g., walking, biking, driving) collected between 2007–2012, resulting in 17,621 trajectories mostly centered around Beijing, China. Porto was originally released as part of the Taxi Service Trajectory challenge and contains data on 442 taxis in Porto, Portugal, spanning more than a year. Cabspotting is also a taxi dataset, with approximately 500 taxis from the San Francisco Bay Area in California, United States, collected over 30 days. For all three datasets, we use time, longitude and latitude. The data is also resampled to one observation per second and map matching is used to reduce observation noise. All models, including baselines and TDDM, are trained and evaluated using the same preprocessed datasets. See C for more details on preprocessing.

4.1 LARGE-SCALE UNCONDITIONAL TRAJECTORY GENERATION

The task of unconditional trajectory generation is to learn a generative model which generate a high quality (I)–(V) synthetic dataset of trajectories. For each dataset, we train each model from scratch and then sample them to generate a synthetic dataset used for evaluation. In the case of TDDM, that means that a region and spatial prior components of the mixture are learned during training, with both relative probability between regions, $p(r_c)$, and spatial priors H_{r_c} for each region partitioned on a grid covering the city in question. The sampling process is then to first draw a region, then draw a trajectory from TDDM conditioned on that region’s spatial prior. Table 1 show the average performance across the three datasets.

TDDM achieves the strongest overall performance, reducing distributional divergences by a large margin compared to all baselines (KL_{sym} : 0.277 vs. 1.153 for Diffusion-TS and 1.232 for DiffTraj, JS : 0.059 vs. 0.198 and 0.209 respectively). This indicates both improved coverage and better avoidance of unrealistic samples. Proportionality measures follow the same trend: TDDM attains the lowest Density and Trip errors (0.019, 0.031), whereas the next-best diffusion models remain higher

(0.029, 0.041). Fidelity is also strong: TSTR reaches 0.011, slightly improving over DiffTraj (0.013) and Diffusion-TS (0.014). On Length error, TDDM matches Diffusion-TS at state of the art (0.004 vs. 0.003), accurately capturing the distribution of distances between consecutive trajectory points (see Appendix F.2 for detailed visualizations), while all other methods trail substantially. Finally, Pattern score is highest for TDDM (0.917), confirming that it best preserves global structure. Overall, TDDM consistently surpasses existing GAN-, VAE-, and diffusion-based methods, combining high fidelity with superior distributional alignment. Results for individual cities is found in the appendix (Table 7).

Inspecting the trajectory samples visually (e.g. Figure 2) we see the advantage of TDDM over the baselines. TDDM generates the trajectories that are the most similar to the original: roads are clearly defined and there are holes in the support where no trajectories appear. Diffusion-TS, DiffTraj and TimeVAE are all capable of capturing the overall shape of the cities, but fail at generating all roads and often generates trajectories that are far from any road. TimeGAN and COSCI-GAN often fail to capture the overall chape of the cities, and TimeGAN especially struggles with mode-collapse. These observations are consistent across all three datasets, see Appendix F.1 for Geolife and Cabsplotting.

4.2 ABLATION STUDY

To probe the role of spatial priors and region size, we conduct ablations (Table 2). Removing spatial priors leaves TSTR unchanged but degrades KL-based scores by up to 5 times, showing that temporal dynamics alone provide useful signals but fail to ensure coverage and proportionality. Reducing partition size to 1×1 km slightly increases Pattern (0.930 vs. 0.917) but worsens Length error (0.150 vs. 0.004), revealing a tradeoff between local coherence and global realism. The full 3×3 km prior-based model provides the best overall balance. Results for individual cities is found in the appendix (Table 8).

Map matching is used as part of the preprocessing to align the raw GPS trajectories to maps for the respective cities, before GPS noise is added back. The same preprocessed data is used throughout the experiments, both for training models and as the target distribution for the evaluation. To verify the effect of map-matching, the top-three models (TDDM, Diffusion-TS, DiffTraj) are also trained without map matching. All models show a significant drop in performance when not using map-matched data, especially for Cabsplotting due to its lower frequency of position updates. The results, shown in Table 9 in the Appendix is consistent with the results where map-matching is used. This demonstrates that TDDM’s improvements in performance compared to the baselines stem from the deaggregation framework.

4.3 OUT-OF-DISTRIBUTION GENERALIZATION

Finally, we evaluate zero-shot generalization (Table 3) for *intra-city* and *city-to-city* transfer. In both settings, models are trained only on source regions and generate for target regions using solely the spatial prior H , with no gradient updates on target trajectories. Results for individual cities, including

Table 1: Evaluation of different models’ performance across several datasets and measures. Models are trained, sampled and evaluated once per dataset. The results are then averaged across datasets.

Measure	TimeGAN	TimeVAE	COSCI-GAN	Diffusion-TS	DiffTraj	TDDM
TSTR (\downarrow)	0.037 ± 0.027	0.018 ± 0.010	0.023 ± 0.007	0.014 ± 0.009	0.013 ± 0.005	0.011 ± 0.006
$KL(S \parallel R)$ (\downarrow)	3.702	2.363	3.046	<u>1.395</u>	1.594	0.301
$KL(R \parallel S)$ (\downarrow)	2.586	1.268	1.740	0.911	<u>0.869</u>	0.253
KL_{sym} (\downarrow)	3.144	1.816	2.393	<u>1.153</u>	1.232	0.277
JS (\downarrow)	0.397	0.287	0.363	<u>0.198</u>	0.209	0.059
KL_{speed} (\downarrow)	0.465	0.225	6.463	<u>0.035</u>	0.126	0.013
Density (\downarrow)	0.258	0.043	0.134	<u>0.029</u>	0.033	0.019
Trip (\downarrow)	0.323	0.056	0.158	<u>0.041</u>	0.042	0.031
Length (\downarrow)	0.097	0.042	0.789	0.003	0.065	<u>0.004</u>
Pattern (\uparrow)	0.677	0.840	0.770	<u>0.907</u>	0.893	0.917

432
 433 Table 2: Ablation study. Models are trained, sampled and evaluated once per dataset. The results are
 434 then averaged across datasets. Note that TDDM has a 3×3 km region size.

Measure	TDDM	1x1 km	w/o spatial prior	w/o spatial prior + rejection
TSTR (\downarrow)	0.011 \pm 0.006	0.024 \pm 0.012	0.011 \pm 0.006	0.014 \pm 0.010
KL($S \parallel R$) (\downarrow)	0.301	<u>0.339</u>	1.569	1.925
KL($R \parallel S$) (\downarrow)	0.253	<u>0.318</u>	1.098	1.252
KL _{sym} (\downarrow)	0.277	<u>0.328</u>	1.334	1.588
JS (\downarrow)	0.059	<u>0.071</u>	0.228	0.266
KL _{speed} (\downarrow)	0.013	0.583	<u>0.323</u>	0.422
Density (\downarrow)	0.019	<u>0.022</u>	0.067	0.063
Trip (\downarrow)	0.031	<u>0.044</u>	0.074	0.081
Length (\downarrow)	0.004	0.150	0.078	<u>0.075</u>
Pattern (\uparrow)	<u>0.917</u>	0.930	0.833	0.860

448 the full city-to-city transfer table is found in the appendix (Table 12). Visualization of the marginal
 449 distribution of the synthetic datasets are shown in Appendix (Figure 15).

450 Zero-shot intra-city transfer, where the model is trained from scratch on only 25% of a map (specific-
 451 ally, a geographically contiguous quadrant as shown in Figures 12–14) and applied to the rest,
 452 shows that aggregated TSTR remains comparable to full coverage (0.010 vs. 0.010), while KL_{sym}
 453 and JS divergences increase (0.545 vs. 0.278; 0.106 vs. 0.059). Pattern also remains high (0.927 vs.
 454 0.940), indicating that spatial priors act as a strong regularizer. Although these aggregated results sug-
 455 gest comparable performance, per-city analysis in the Appendix reveals more substantial variations
 456 (Table 12).

457 In zero-shot city-to-city transfer, performance varies with the source dataset but remains competitive:
 458 Pattern stays above 0.915 across all cases, proportionality measures are consistently lower than
 459 GAN/VAEs and close to diffusion baselines, and TSTR often matches in-distribution performance
 460 when trained on Porto. The main weakness is Length error, which increases to 0.06–0.11 across cities,
 461 suggesting that fine-grained distance modeling is less transferable. Despite this, TDDM demon-
 462 strates robust fidelity and distributional generalization across cities, highlighting the effectiveness of
 463 separating spatial priors from temporal dynamics.

464 Interestingly, cross-city transfer from Porto often yields stronger results than training on limited
 465 portions of the target city. On average, models trained on Porto generalize with lower KL and JS
 466 divergences (0.335 and 0.071) than those trained on only 25% of the target city (0.545 and 0.106),
 467 and they also maintain slightly better proportionality and Pattern scores (0.930 vs. 0.927). The
 468 only exception is Length error, where access to even a small fraction of local data provides an
 469 advantage (0.026 vs. 0.060). This reflects city-specific differences in trajectory length distributions,
 470 which cannot be inferred from spatial priors alone. For instance, Porto exhibits a heavier-tailed
 471 distribution than Cabspotting (Appendix F.2, Figure 19). This suggests that Porto captures temporal
 472 dynamics and spatial statistics that are broadly representative across cities, making it an unexpectedly
 473 strong *universal source* dataset for this setting. More generally, these results highlight a tradeoff:
 474 if path-length accuracy is paramount, local data (even in small amounts) remains valuable, but for
 475 distributional coverage and spatial structure, carefully chosen training cities may outperform partial
 476 local coverage.

486
 487 Table 3: Generalization performance across different training scenarios. Left: Intra-city generalization
 488 comparing training on the partial (25%) versus full (100%) spatial domain. Right: City-to-city
 489 generalization where models are trained on one dataset and evaluated on others. Results are aggregated
 490 across datasets.

Measure	Intra-city (Training data)		Geolife	City-to-city (Trained on)	
	25%	100%		Porto	Cabspotting
TSTR (↓)	0.010 ± 0.006	0.010 ± 0.007	0.016 ± 0.008	0.010 ± 0.005	0.011 ± 0.006
KL($S \parallel R$) (↓)	<u>0.615</u>	0.305	0.903	0.357	<u>0.610</u>
KL($R \parallel S$) (↓)	<u>0.474</u>	0.251	0.688	0.313	<u>0.449</u>
KL _{sym} (↓)	<u>0.545</u>	0.278	0.795	0.335	<u>0.530</u>
JS (↓)	<u>0.106</u>	0.059	0.149	0.071	<u>0.102</u>
KL _{speed} (↓)	<u>0.101</u>	0.012	<u>0.322</u>	0.238	0.393
Density (↓)	<u>0.021</u>	0.015	0.018	0.018	<u>0.022</u>
Trip (↓)	<u>0.036</u>	0.027	0.031	<u>0.036</u>	0.042
Length (↓)	<u>0.026</u>	0.003	<u>0.082</u>	0.060	0.109
Pattern (↑)	<u>0.927</u>	0.940	<u>0.925</u>	0.930	0.915

5 CONCLUSION

We have presented the Temporal Deaggregation Diffusion Model (TDDM), a hierarchical generative framework that separates spatial priors from temporal dynamics for large-scale trajectory generation. Across three major urban datasets, TDDM consistently improves distributional alignment, achieving up to 4 times lower KL divergences than the best diffusion baselines, while also setting new state of the art on Density, Trip, and Pattern measures. Importantly, TDDM matches leading models on Length error and outperforms them on fidelity as measured by TSTR. This is also confirmed visually.

Beyond unconditional generation, ablation studies highlight the critical role of spatial priors for coverage and proportionality, while generalization experiments demonstrate that TDDM can synthesize realistic trajectories in unseen parts of cities and across entirely new cities. In particular, we find that training on Porto generalizes better on average to other cities than training on partial local data, suggesting that certain cities may act as representative source datasets. That is, datasets that generalize broadly across urban contexts and provide stronger transferability than limited amounts of local data.

Taken together, these results show that factorizing *where* and *how* people move not only advances generation quality but also unlocks strong out-of-distribution generalization. TDDM thus provides both a methodological advance and a practical step toward scalable and generalizable mobility modeling.

Future Work. Several promising directions remain for extending TDDM. First, augmenting the spatial prior H with additional marginal information (such as trajectory length distributions, time-of-day priors, or directional priors for traffic handedness) could improve temporal fidelity and enable accurate generalization to left-hand vs. right-hand traffic patterns. Second, incorporating road hierarchy information as priors would better capture local speed structure and road network topology. Finally, extending H with city-specific temporal marginals would likely improve cross-city temporal fidelity while preserving the zero-shot spatial generalization capabilities demonstrated in this work.

540
541
ETHICS STATEMENT

542 This work focuses on improving the quality and generalization capabilities of synthetic human
 543 mobility trajectory generation using publicly available GPS datasets (Geolife, Porto, Cabspotting)
 544 with appropriate licenses as detailed in Appendix B.1. We acknowledge that synthetic trajectory
 545 generation technology has dual-use potential. While our primary motivation is to advance the state-of-
 546 the-art in generative modeling to enable beneficial applications such as urban planning, transportation
 547 research, and mobility simulation, we recognize that high-quality synthetic data generation techniques
 548 could potentially be misused for surveillance or other harmful purposes. We encourage responsible
 549 use of our methods and emphasize that practitioners should consider appropriate safeguards when
 550 deploying synthetic trajectory generation in real-world applications. Additionally, we note that
 551 privacy considerations must be carefully evaluated when applying our methods, as even publicly
 552 available trajectory datasets can potentially be used for re-identification when combined with other
 553 data sources or when analyzed with sufficient temporal and spatial resolution. We emphasize
 554 that our methods do not guarantee privacy-preserving synthetic data generation, and we strongly
 555 encourage practitioners to combine our approach with established privacy-preserving techniques such
 556 as differential privacy when working with real-world trajectory data. Our research contributes to
 557 the broader goal of developing high-fidelity generative models that can support legitimate research
 558 and planning activities while reducing the need for access to sensitive real trajectory data. We
 559 have designed our evaluation framework to focus on distributional properties and generation quality
 560 without enabling inference about specific individuals in the original datasets.

561
562
REPRODUCIBILITY

563 To ensure full reproducibility of our results, we provide comprehensive implementation details and
 564 experimental resources. We release complete, runnable source code for all experiments ¹, including
 565 implementations of TDDM and all baseline methods, along with evaluation code for all measures
 566 used. The core TDDM algorithm is detailed in Algorithms 1 and 2, with complete architectural
 567 specifications in Section 3 and extended implementation details in Appendix C. All hyperparameters
 568 used in our experiments are provided in Appendix C.4, including the search space and final values
 569 obtained through systematic optimization. Dataset preprocessing steps for all three datasets (Geolife,
 570 Porto, Cabspotting) are comprehensively described in Appendix B, including data licenses, filtering
 571 criteria, and map-matching procedures. Our evaluation methodology is detailed in Section 4 with
 572 extended measure descriptions in Appendix E. All experimental configurations, including baseline
 573 implementations and evaluation protocols, are provided to enable direct replication of our results
 574 across all datasets and experimental conditions.

575
576
LLM USAGE

577 This paper was developed with assistance from large language models as writing and research tools,
 578 including Claude (Anthropic), ChatGPT (OpenAI), and Cursor AI for code assistance. The LLMs
 579 contributed to:

- 580 • **Writing refinement:** Improving clarity, flow, and exposition throughout the paper
- 581 • **Technical positioning:** Helping articulate the methodological approach and frame the
 582 contribution within existing literature
- 583 • **Experimental analysis:** Assisting with evaluation design and results presentation
- 584 • **Literature review:** Supporting the discovery and organization of related work
- 585 • **Code development:** Assisting with implementation and debugging of code

586 All research concepts, experimental design, implementation, and scientific conclusions are the original
 587 work of the authors. All LLM-generated content has been verified and validated by the authors. The
 588 authors take full responsibility for all content and claims presented in this paper.

589
590 ¹Available at <https://anonymous.4open.science/r/tddm/>.

594 REFERENCES
595

596 Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna:
597 A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM
598 SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.

599 Ahmed Alaa, Boris Van Breugel, Evgeny S. Saveliev, and Mihaela van der Schaar. How faithful
600 is your synthetic data? Sample-level metrics for evaluating and auditing generative models. In
601 Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato
602 (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of
603 *Proceedings of Machine Learning Research*, pp. 290–306. PMLR, 17–23 Jul 2022.

604 Juan Lopez Alcaraz and Nils Strothoff. Diffusion-based time series imputation and forecasting
605 with structured state space models. *Transactions on Machine Learning Research*, 2023. ISSN
606 2835-8856.

607 608 Alfredo Alessandrini, Daniela Ghio, Silvia Migali, et al. *Estimating net migration at high spatial
609 resolution*. Publications Office of the European Union, 2020.

610 Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen,
611 Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al.
612 Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.

613 614 Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky,
615 Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will
616 Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael
617 Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos,
618 Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, CK Luk, Bert Maher, Yunjie Pan, Christian
619 Puhrsch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Michael Suo, Phil
620 Tillet, Eikan Wang, Xiaodong Wang, William Wen, Shunting Zhang, Xu Zhao, Keren Zhou,
621 Richard Zou, Ajit Mathews, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster
622 Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation.
623 In *29th ACM International Conference on Architectural Support for Programming Languages and
Operating Systems, Volume 2 (ASPLOS '24)*. ACM, April 2024. doi: 10.1145/3620665.3640366.

624 Erik Buchholz, Alsharif Abuadbba, Shuo Wang, Surya Nepal, and Salil S Kanhere. Sok: Can
625 trajectory generation combine privacy and utility? *arXiv preprint arXiv:2403.07218*, 2024.

626 627 Chu Cao and Mo Li. Generating mobility trajectories with retained data utility. In *Proceedings of the
27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 2610–2620, 2021.

628 629 Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference
on machine learning*, pp. 2990–2999. PMLR, 2016.

630 632 Yun Dai, Chao Yang, Kaixin Liu, Angpeng Liu, and Yi Liu. Timedddpm: Time series augmentation
633 strategy for industrial soft sensing. *IEEE Sensors Journal*, 2023.

634 635 Abhyuday Desai, Cynthia Freeman, Ian Beaver, and Zuhui Wang. TimeVAE: A variational auto-
encoder for multivariate time series generation. *arXiv preprint arXiv:2111.08095*, 2021.

636 637 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
638 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
639 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint
arXiv:2010.11929*, 2020.

640 641 Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. Real-valued (Medical) Time Series
642 Generation with Recurrent Conditional GANs. *arXiv preprint arXiv:1706.02633*, 2017.

643 644 Shibo Feng, Chunyan Miao, Zhong Zhang, and Peilin Zhao. Latent diffusion transformer for proba-
645 bility time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
646 volume 38, pp. 11979–11987, 2024.

647 Colin Goodall. Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical
Society: Series B (Methodological)*, 53(2):285–339, 1991.

648 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Advances*
 649 *in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc.,
 650 2020.

651

652 Cornelia Ilin, Sébastien Annan-Phan, Xiao Hui Tai, Shikhar Mehra, Solomon Hsiang, and Joshua E
 653 Blumenstock. Public mobility data enables covid-19 forecasting and management at local and
 654 global scales. *Scientific reports*, 11(1):13531, 2021.

655

656 Jinsung Jeon, Jeonghak Kim, Haryong Song, Seunghyeon Cho, and Noseong Park. Gt-gan: General
 657 purpose time series synthesis with generative adversarial networks. *Advances in Neural Information*
 658 *Processing Systems*, 35:36999–37010, 2022.

659

660 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
 661 *arXiv:1412.6980*, 2014.

662

663 Ibai Lana, Javier Del Ser, Manuel Velez, and Eleni I Vlahogianni. Road traffic forecasting: Recent
 664 advances and new challenges. *IEEE Intelligent Transportation Systems Magazine*, 10(2):93–109,
 2018.

665

666 Luis Moreira-Matias, Michel Ferreira, Joao Mendes-Moreira, L. L., and J. J. Taxi Service Trajectory
 667 - Prediction Challenge, ECML PKDD 2015. UCI Machine Learning Repository, 2013. DOI:
 668 <https://doi.org/10.24432/C55W25>.

669

670 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.
 671 In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on*
 672 *Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8162–8171.
 673 PMLR, 18–24 Jul 2021.

674

675 Venla Niva, Alexander Horton, Vili Virkki, Matias Heino, Maria Kosonen, Marko Kallio, Pekka
 676 Kinnunen, Guy J Abel, Raya Muttarak, Maija Taka, et al. World’s human migration patterns in
 677 2000–2019 unveiled by high-resolution data. *Nature Human Behaviour*, 7(11):2023–2037, 2023.

678

679 Michal Piorkowski, Natasa Sarafijanovic-Djukic, and Matthias Grossglauser. Crawdad epfl/mobility,
 680 2022.

681

682 Ali Seyfi, Jean-Francois Rajotte, and Raymond Ng. Generating multivariate time series with CCommon
 683 Source CoordInated GAN (COSCI-GAN). *Advances in Neural Information Processing Systems*,
 684 35:32777–32788, 2022.

685

686 Lifeng Shen and James Kwok. Non-autoregressive conditional diffusion models for time series
 687 prediction. In *International Conference on Machine Learning*, pp. 31016–31029. PMLR, 2023.

688

689 Xiao Hui Tai, Shikhar Mehra, and Joshua E Blumenstock. Mobile phone data reveal the effects of
 690 violence on internal displacement in afghanistan. *Nature human behaviour*, 6(5):624–634, 2022.

691

692 Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. CSDI: Conditional score-based diffu-
 693 sion models for probabilistic time series imputation. In M. Ranzato, A. Beygelzimer, Y. Dauphin,
 694 P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*,
 695 volume 34, pp. 24804–24816. Curran Associates, Inc., 2021.

696

697 Valhalla Contributors. Valhalla: Open source routing engine for openstreetmap, 2025. URL <https://github.com/valhalla/valhalla>. MIT License.

698

699 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 700 Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. Von Luxburg, S. Bengio,
 701 H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information*
 702 *Processing Systems*, volume 30, 2017.

703

704 Huandong Wang, Changzheng Gao, Yuchen Wu, Depeng Jin, Lina Yao, and Yong Li. Pategail: a
 705 privacy-preserving mobility trajectory generator with imitation learning. In *Proceedings of the*
 706 *AAAI Conference on Artificial Intelligence*, volume 37, pp. 14539–14547, 2023.

702 Yu Wang, Tongya Zheng, Yuxuan Liang, Shunyu Liu, and Mingli Song. Cola: Cross-city mobility
 703 transformer for human trajectory simulation. In *Proceedings of the ACM on Web Conference 2024*,
 704 pp. 3509–3520, 2024.

705 Yun Wang, Faiz Currim, and Sudha Ram. Deep learning of spatiotemporal patterns for urban mobility
 706 prediction using big data. *Information Systems Research*, 33(2):579–598, 2022.

708 Qitian Wu, Rui Gao, and Hongyuan Zha. Bridging explicit and implicit deep generative models via
 709 neural stein estimators. *Advances in Neural Information Processing Systems*, 34:11274–11286,
 710 2021.

711 Gang Xiong, Zhishuai Li, Meihua Zhao, Yu Zhang, Qinghai Miao, Yisheng Lv, and Fei-Yue Wang.
 712 Trajsgan: A semantic-guiding adversarial network for rban trajectory generation. *IEEE Trans-
 713 actions on Computational Social Systems*, 11(2):1733–1743, 2024. doi: 10.1109/TCSS.2023.
 714 3235923.

716 Jinsung Yoon, Daniel Jarrett, and Mihaela Van der Schaar. Time-series Generative Adversarial
 717 Networks. *Advances in neural information processing systems (NeurIPS)*, 32, 2019a.

718 Jinsung Yoon, James Jordon, and Mihaela van der Schaar. PATE-GAN: Generating synthetic data
 719 with differential privacy guarantees. In *International Conference on Learning Representations*,
 720 2019b.

722 Xinyu Yuan and Yan Qiao. Diffusion-TS: Interpretable diffusion for general time series generation.
 723 In *The Twelfth International Conference on Learning Representations*, 2024.

724 Yu Zheng, Hao Fu, Xing Xie, Wei-Ying Ma, and Quannan Li. *Geolife GPS trajectory dataset - User
 725 Guide*, geolife gps trajectories 1.1 edition, July 2011.

727 Yuanshao Zhu, Yongchao Ye, Shiyao Zhang, Xiangyu Zhao, and James Yu. Difftraj: Generating gps
 728 trajectory with diffusion probabilistic model. *Advances in Neural Information Processing Systems*,
 729 36:65168–65188, 2023.

730 Yuanshao Zhu, James Jianqiao Yu, Xiangyu Zhao, Qidong Liu, Yongchao Ye, Wei Chen, Zijian
 731 Zhang, Xuetao Wei, and Yuxuan Liang. Controltraj: Controllable trajectory generation with
 732 topology-constrained diffusion model. In *Proceedings of the 30th ACM SIGKDD Conference on
 733 Knowledge Discovery and Data Mining*, pp. 4676–4687, 2024.

734
 735
 736
 737
 738
 739
 740
 741
 742
 743
 744
 745
 746
 747
 748
 749
 750
 751
 752
 753
 754
 755

APPENDIX: SUPPLEMENTARY MATERIAL FOR LEARNING TO DEAGGREGATE

The appendix contains supplementary material to the paper *Learning to Deaggregate: Unconditional Generation of Trajectories at Scale*. We provide detailed information on related work, model architecture, implementation details, extended experimental results, and additional visualizations.

Contents

A. Extended Related Work	15
B. Datasets and Preprocessing	17
B.1. Dataset Descriptions and Licenses	17
B.2. Preprocessing Pipeline	17
C. Architecture Details	18
C.1. Denoising Diffusion Probabilistic Models	18
C.2. Optimization Details	18
C.3. Encoder Input	18
C.4. Hyperparameters	19
D. Computational Complexity and Runtime Analysis	20
D.1. Theoretical Complexity	20
D.2. Empirical Runtime and Memory Requirements	20
E. Extended Evaluation Measure Description	21
F. Additional Results	23
F.1. Additional Result Visualizations	23
F.3. Quantitative Results Tables	40

A EXTENDED RELATED WORK

In Section 1 and 4 of the main paper, we described the state-of-the-art models used for comparison. Here, we provide extended details on how these models differ from our proposed TDDM approach, particularly regarding their conditional generation capabilities and architectural choices.

Previous work on unconditional generation of time-series data has focused on variations of the generative adversarial networks (GAN) architecture (Esteban et al., 2017; Yoon et al., 2019a; Jeon et al., 2022) and, more recently diffusion models (Zhu et al., 2023; Yuan & Qiao, 2024).

There has also been an interest in using time-series generation for imputation and forecasting (Alcaraz & Strodthoff, 2023; Tashiro et al., 2021; Shen & Kwok, 2023; Dai et al., 2023; Feng et al., 2024). Transformer-based time-series foundation models have been proposed as a general purpose forecasting tool (Ansari et al., 2024), but has not been evaluated on the unconditional generation task.

TimeGAN (Yoon et al., 2019a) consists of a generative adversarial network (GAN) operating inside the latent space of an autoencoder. To further improve performance, they add an additional network with the task of predicting one time step ahead. The encoder, decoder, supervisor, generator and discriminator are all implemented using autoregressive models and in practice they use gated recurrent units (GRUs).

The TimeVAE (Desai et al., 2021) architecture is a variant of the popular variational autoencoder architecture. The autoencoder is trained with an additional loss component to have the latent space conform to a known statistical distribution, in this instance a multivariate normal distribution. The

810 autoencoder is trained to both minimize the reconstruction loss, as well as minimizing the divergence
 811 between the embedded data and the prior set for the latent space.

812 COSCI-GAN (Seyfi et al., 2022) proposes to use a separate generative adversarial network for each
 813 channel of the data. The individual GANs all share single source of noise as input to the generator
 814 and, additionally they have a central discriminator that is given the stacked output from the all the
 815 generators as input.

816 Diffusion-TS (Yuan & Qiao, 2024) adapts the denoising diffusion architecture (Ho et al., 2020)
 817 to generate time-series data by implementing the denoising step with a multilayer neural network
 818 each consisting of a transformer block, a fully connected neural network as well as time-series
 819 specific layers with the aim of improved performance and interpretability. The architecture also has
 820 support for conditional generation to enable using the models for imputation and forecasting. As
 821 Diffusion-TS moves away from the U-Net approach used in trajectory-specific models to a partial
 822 Transformer architecture, we extend this further to a full Transformer-based architecture, avoiding
 823 additional time-series specific induction biases while gaining increased performance for synthetic
 824 data generation.

825 Recent work has extended these generative architectures to the specific domain of GPS trajectory
 826 generation, with varying approaches to conditional control.

827 DiffTraj (Zhu et al., 2023) adapts the denoising diffusion architecture (Ho et al., 2020) to generate
 828 GPS-trajectories. The U-Net architecture used in (Ho et al., 2020) for denoising, is also adapted to
 829 work with trajectories. As part of the generation process, the model is conditioned on individual
 830 trajectory-specific statistics such as velocity, distance, and departure time.

831 ControlTraj (Zhu et al., 2024) is a diffusion-based framework for generating GPS trajectories conditioned
 832 on road network topology from OpenStreetMap and trip-specific attributes (departure time,
 833 distance, speed). Similar to DiffTraj, ControlTraj uses trajectory-level conditioning to control individual
 834 samples. In contrast, TDDM conditions on spatial priors (aggregate marginal occupancy
 835 distributions computed from multiple trajectories).

836 COLA (Wang et al., 2024) is a model for generating discrete trajectories, i.e., where space is
 837 discretized into location IDs, and where samples are temporally sparser (hourly timesteps vs. 1-second
 838 for TDDM). Both COLA and TDDM address city-to-city transfer, but using different paradigms.
 839 COLA uses transfer learning with gradient-based adaptation, while TDDM achieves zero-shot transfer
 840 via spatial prior conditioning. The zero-shot approach only requires marginal occupancy information,
 841 while the transfer learning approach requires access to full trajectories from the target domain.

842 TrajGen (Cao & Li, 2021) is an unconditional generative model for trajectories. It uses an unconditional
 843 generative image generative model to generate images of trajectories. The images show a
 844 top-down view of single trajectories. These are then decoded into a sequence of points, which are
 845 then ordered and then timestamped. While both TrajGen and TDDM separate spatial and temporal
 846 information, they differ fundamentally in representation and conditioning: TrajGen’s unconditional
 847 approach means that once trained, the model has limited control of what trajectory patterns are
 848 generated, whereas TDDM explicitly conditions on spatial priors H , enabling controllable generation
 849 and zero-shot cross-city transfer.

850 While DiffTraj, ControlTraj, and Diffusion-TS all support conditional generation, our approach
 851 differs fundamentally in the type and purpose of conditioning information. DiffTraj and ControlTraj
 852 condition on sample-specific statistics (velocity, distance, road networks), creating a near one-to-
 853 one mapping between conditions and individual trajectories. Diffusion-TS uses conditioning for
 854 enabling imputation and forecasting rather than improving generation quality. In contrast, TDDM
 855 conditions on spatial priors (aggregate marginal occupancy distributions computed from multiple
 856 trajectories), which cannot be directly tied to any individual sample. This distinction enables TDDM
 857 to generalize to entirely new environments through zero-shot transfer (requiring only marginal
 858 occupancy information), while sample-specific conditioning limits generation to the trained spatial
 859 region and increases memorization risk. TrajGen’s unconditional approach offers no control over
 860 generated patterns, while COLA’s transfer learning paradigm requires access to full trajectories from
 861 the target domain for gradient-based adaptation.

862
 863

864 **B DATASETS AND PREPROCESSING**
865866 In this section, we describe the datasets used in our experiments and the preprocessing steps applied
867 to prepare the data for training and evaluation.
868869 The datasets described in Section 4 of the main paper provide geographical diversity across different
870 continents while ensuring sufficient data density for meaningful evaluation. Our preprocessing
871 pipeline, described below, is applied uniformly to all methods to ensure fair comparison.
872873 The trajectory segmentation process helps all methods by creating cleaner, more consistent training
874 data with fewer anomalies. It is important to note that while our model operates on such local
875 regions during both training and inference, we do not generate trajectories that extend beyond the
876 size of regions used during training. Instead, the hierarchical approach allows us to generate coherent
877 trajectories for multiple adjacent regions that together compose large-scale environments.
878879 **B.1 DATASET DESCRIPTIONS AND LICENSES**
880881 All three datasets used in this work are publicly available and contain GPS trajectory data from
882 different geographical regions:
883884

- **Geolife** (Zheng et al., 2011): Data from 182 users collected between 2007-2012, resulting
885 in 17,621 trajectories mostly centered around Beijing, China. Released under the Microsoft
886 Research License Agreement, which allows academic research but prohibits redistribution.
887
- **Porto** (Moreira-Matias et al., 2013): Released as part of the Taxi Service Trajectory chal-
888 lenge, containing data on 442 taxis in Porto, Portugal, spanning more than a year. Released
889 under the CC BY 4.0 license.
890
- **Cabspotting** (Piorkowski et al., 2022): Taxi dataset with approximately 500 taxis from the
891 San Francisco Bay Area in California, United States, collected over 30 days. Released under
892 the CC BY 4.0 license. An account is required to download this data.
893

894 For all three datasets, we use time, longitude and latitude as the primary features.
895896 **B.2 PREPROCESSING**
897898 Raw GPS data contains multiple problems:
899900

- GPS drift
- GPS being turned on even when the vehicle is not in use
- large spikes in velocity

901 We split trajectories into sub-trajectories if any observation leaves the geographic bounds, exceeds
902 the velocity limit or if too long time has passed since the previous observation. In the raw data, a
903 trajectory can span several days with several hours between observations. Introducing a time-limit
904 allows us to break these trajectories into individual and more time-constrained journeys.
905906 We up-sample the data to one observation per second. Map matching using Valhalla (Valhalla
907 Contributors, 2025) is then used to map the GPS trajectories to the road network, the distance
908 between the matched and interpolated trajectories are used to estimate a noise distribution. The final
909 training data consists of the map matched trajectories with trajectory-noise added from the noise
910 distribution, i.e. we sample the noise distribution once per trajectory.
911912 **B.3 HARDWARE**
913914 All experiments were conducted on a system with the following specifications:
915916

- **CPU:** AMD Ryzen 9 5900X
- **RAM:** 128 GB
- **GPU:** NVIDIA GeForce RTX 3090 Ti (24 GB GDDR6X)

918 **C ARCHITECTURE DETAILS**
 919

920 **C.1 DENOISING DIFFUSION PROBABILISTIC MODELS**
 921

922 Denoising diffusion probabilistic includes two processes: a known noise-adding process and an
 923 approximated denoising process. The noise-adding process starts at a random sample $x \in \mathbb{X}$,
 924 denoted x_0 where 0 is the noise step, and adds noise at each steps according to a noise schedule
 925 $\beta = \{\beta_1, \dots, \beta_T\}$. By using a Gaussian noise process (Ho et al., 2020), the forward process has
 926 a closed expression for time point t : $q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)I)$, where $\alpha_t := 1 - \beta_t$,
 927 $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$. The reverse process is starts at $p(x_T) = \mathcal{N}(x_T; 0, I)$ and the transition function for
 928 $1 < t \leq T$ is learned:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (6)$$

929 There are different ways to parameterize $\mu_\theta(x_t, t)$ and $\Sigma_\theta(x_t, t)$ and we use the noise prediction
 930 parameterization (Ho et al., 2020):
 931

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right), \quad \Sigma_\theta(x_t, t) = \sigma_t^2 I \quad (7)$$

932 where $\epsilon_\theta(x_t, t)$ is a learned function and is trained to predict the added noise at step t of the noise
 933 adding process. In this work, we extend ϵ_θ with discretized marginal distribution H as described in
 934 the main paper in Section 3.
 935

936 **C.2 OPTIMIZATION DETAILS**
 937

938 We use the following positional encoding to encode a one-dimensional signal (Vaswani et al., 2017):
 939

$$\text{PE}_{(pos, 2i)} = \sin \left(-e^{i \frac{\log(10000)}{\frac{d}{2} - 1}} \right) \quad (8)$$

$$\text{PE}_{(pos, 2i+1)} = \cos \left(-e^{i \frac{\log(10000)}{\frac{d}{2} - 1}} \right), \text{ where} \quad (9)$$

940 pos is the signal, e.g. a position x , position y or time point t , and i is the target dimension. For
 941 position, we encode a single position observation in \mathbb{R} into \mathbb{R}^D , so the function would be called once
 942 for $i \in 1 \dots D$. We use the cosine noise schedule as proposed by Nichol & Dhariwal (2021):
 943

$$\bar{\alpha}_t = \frac{f(t)}{f(0)}, f(t) = \cos \left(\frac{t/T + s}{1 + s} \cdot \frac{\pi}{2} \right)^2 \quad (10)$$

944 with $s = 0.008$. Furthermore, we use the simplified loss function proposed by Ho et al. (2020):
 945

$$\mathcal{L}_{\text{simple}}(\theta) := \mathbb{E}_{t, x_0, \epsilon} \left[\left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t) \right\|^2 \right] \quad (11)$$

946 We use the Adam (Kingma & Ba, 2014) optimizer and implement our model in PyTorch (Ansel et al.,
 947 2024).
 948

949 **C.3 ENCODER INPUT**
 950

951 We set N to be a hyperparameter, and D is the number of dimensions of a single observation in a
 952 trajectory, in our case $D = 2$. The input to the transformer encoder is:
 953

- 954 • L input tokens, each token corresponding to a time point in the noisy sequence. **Note:** The
 955 token size depends on the number of features of the dataset. It is a concatenation of:
 - 956 – $x \in \mathbb{R}^{DN}$, each corresponding to a dimension observed at each time point encoded
 957 using positional encoding
 - 958 – $x \in \mathbb{R}^N$, the time point encoded using the positional encoding introduced in (Vaswani
 959 et al., 2017)

972 – $x \in \mathbb{R}^N$, a learned vector encoding denoting that this is a token that corresponds to a
 973 noisy sequence
 974
 975 • **Conditional information:** 64 tokens, each corresponding to a patch of the heatmap and
 976 being a concatenation of:
 977 – $x \in \mathbb{R}^N$, corresponding to the x position of the patch. Encoded using positional
 978 encoding (Vaswani et al., 2017).
 979 – $x \in \mathbb{R}^N$, corresponding to the y position of the patch. Encoded using positional
 980 encoding (Vaswani et al., 2017).
 981 – $x \in \mathbb{R}^N$, corresponding to the intensity of the heatmap. Encoded using a linear
 982 projection (Dosovitskiy et al., 2020).
 983 – $x \in \mathbb{R}^N$, a learned vector encoding denoting that this is a token that corresponds to the
 984 conditional information
 985
 986 • A token encoding the current denoising step:
 987 – $x \in \mathbb{R}^{N(D+1)}$, the denoising step encoded using positional encoding (Vaswani et al.,
 988 2017)
 989 – $x \in \mathbb{R}^N$, a learned vector encoding denoting that this is a token that corresponds to the
 990 denoising step
 991
 992

993 C.4 HYPERPARAMETERS

995 We employ Optuna (Akiba et al., 2019) with Tree-structured Parzen Estimator (TPE) sampling to
 996 optimize our model’s hyperparameters. The optimization process consists of 50 trials, where the
 997 model is trained from scratch trains for 100 epochs. For each trial, we sample hyperparameters from
 998 the defined search space and evaluate the model’s performance by calculating the Jensen-Shannon
 999 divergence between the marginal distributions of the synthetic data and the marginal distributions of
 1000 the training data. We only use the Geolife dataset for hyperparameter tuning. The hyperparameters
 1001 included in the optimization are shown in Table 4 and the final values are shown in Table 5.

1002
 1003 Table 4: Hyperparameter Search Space

1005 Parameter	1006 Search Space	1007 Description
1006 Hidden Dimension	1007 $\{4, 8, 16, 32, 64, 128\}$	1008 Size of transformer hidden layers
1008 Number of Layers	1009 $\{1, 2, 4, 8\}$	1010 Depth of transformer architecture
1009 Number of Attention Heads	1010 $\{1, 2, 4, 8\}$	1011 Multi-head attention mechanism size
1010 Diffusion Timesteps	1011 $\{100, 500, 1000\}$	1012 Number of diffusion steps
1011 Learning Rate	1012 $[10^{-6}, 10^{-2}]$	1013 Log-uniform sampling

1014 The Jensen-Shannon divergence was chosen as the optimization metric because it directly aligns with
 1015 our goal of generating synthetic trajectories that match the distributional properties of real data, as
 1016 discussed in Section 4 where we evaluate models on both sample-level fidelity and distribution-level
 1017 similarity.

1018 Table 5: Final Hyperparameters

1019 Parameter	1020 Value
1021 Hidden Dimension	1022 128
1022 Number of Layers	1023 8
1023 Number of Attention Heads	1024 2
1024 Diffusion Timesteps	1025 500
1025 Learning Rate	0.00017483

1026 D COMPUTATIONAL COMPLEXITY AND RUNTIME ANALYSIS 1027

1028 In this section, we provide a detailed analysis of TDDM’s computational complexity and empirical
1029 runtime characteristics. This addresses concerns about the scalability and practical deployment of our
1030 method for large-scale trajectory generation.
1031

1032 D.1 THEORETICAL COMPLEXITY 1033

1034 The computational complexity of TDDM is determined by the transformer encoder architecture used
1035 for denoising at each diffusion step. For a single denoising step, the per-trajectory cost is:
1036

$$1037 O((L + R)d^2 + (L + R)^2d) \quad (12)$$

1039 where:

- 1040 • L is the trajectory sequence length (number of time steps)
- 1041 • R is the number of spatial prior patches (64 in our implementation)
- 1042 • d is the hidden dimension of the transformer (128 in our final model)

1044 The first term $(L + R)d^2$ corresponds to the cost of the feed-forward layers in the transformer, while
1045 the second term $(L + R)^2d$ corresponds to the cost of the self-attention mechanism that processes
1046 both trajectory tokens and spatial prior tokens jointly.
1047

1048 **Scalability to large cities.** Importantly, because we operate on fixed-size 3×3 km regions with
1049 a constant number of spatial prior patches ($R = 64$), the per-trajectory computational cost remains
1050 *constant* across cities of different sizes. City-scale generation scales *linearly* with geographic coverage
1051 (city area), as we independently process each region and stitch the results together according to the
1052 mixture model formulation in Section 3. This is in contrast to methods that must process entire
1053 city-scale contexts, which would scale quadratically with area in attention-based models.
1054

1055 D.2 EMPIRICAL RUNTIME AND MEMORY REQUIREMENTS

1056 We measure TDDM’s inference performance on an NVIDIA GeForce RTX 3090 Ti (24 GB GDDR6X)
1057 for different batch sizes. Table 6 shows the time per trajectory and peak GPU memory consumption.
1058

1059 Table 6: TDDM inference performance across different batch sizes on RTX 3090 Ti. Time per trajec-
1060 tory decreases with larger batch sizes due to improved GPU utilization, while memory consumption
1061 grows sub-linearly.
1062

1063 Batch Size	1064 Time per Trajectory (ms)	1065 Peak GPU Memory (MB)
1065 1	1162	313
1066 8	427	340
1067 64	349	558
1068 256	340	1307

1069 Key observations from these measurements:
1070

- 1071 • **Batching efficiency:** Increasing batch size from 1 to 256 reduces per-trajectory time by
1072 approximately 3.4x (from 1162 ms to 340 ms), demonstrating good GPU utilization at
1073 higher batch sizes.
- 1074 • **Memory efficiency:** Peak GPU memory grows sub-linearly with batch size. A 256x increase
1075 in batch size only requires approximately 4.2x increase in memory (from 313 MB to 1307
1076 MB), indicating efficient memory usage that can accommodate large-scale generation on
1077 modern GPUs.
- 1078 • **Practical throughput:** At the batch size of 256, TDDM can generate approximately 2.9
1079 trajectories per second on a single RTX 3090 Ti. For a typical city-scale dataset with 50,000
trajectories, complete generation takes approximately 4.7 hours.

- **Comparison with baselines:** While direct runtime comparisons depend on implementation details, hardware, and sampling procedures, TDDM’s per-trajectory inference time of 340ms is competitive with other diffusion-based methods. The ability to batch efficiently and the constant per-trajectory cost (independent of city size) make TDDM practical for large-scale applications.

E EXTENDED EVALUATION MEASURES DESCRIPTION

In Section 4, we introduced evaluation measures spanning multiple quality dimensions for synthetic trajectory data. Here we provide a more detailed description of these measures, their theoretical foundations, and how they complement each other in our benchmark.

E.1 DIFFTRAJ MEASURES

The measures from DiffTraj (Zhu et al., 2023) are²:

- **Density Error:** A pair of heatmaps of the training and synthetic data are calculated by dividing the city into 16×16 blocks. The number of observations in each block are counted and normalized. The Jensen-Shannon divergence is calculated between the training data heatmap and synthetic data heatmap.
- **Trip Error:** Two heatmap pairs, each 16×16 blocks, are calculated. The first pair is of the start positions for all trajectories, one heatmap for the training data and another for the synthetic data. The second pair is calculated from the last position for all trajectories. The Jensen-Shannon divergence is calculated once for each pair and then the average is reported.
- **Length Error:** The distance between consecutive observations are calculated, once for the training data and once for the synthetic data. Histograms are calculated for each, with the number of bins set to 16. Finally, the Jensen-Shannon divergence is calculated between the histogram of training data and the histogram of the synthetic data.
- **Pattern Score:** Using the heatmaps from Density Error, the top N areas (highest count) from the training and synthetic data are collected. The F-score is then calculated and reported.

E.2 KL-BASED DISTRIBUTION MEASURES

Our evaluation framework incorporates several Kullback-Leibler (KL) divergence-based measures that address different quality dimensions of synthetic trajectory data:

- **KL($R \parallel S$):** Kullback-Leibler divergence of the real (R) distribution from the synthetic (S) distribution measures how well the synthetic distribution covers areas where real data exists:

$$KL(R \parallel S) = \sum_{i,j} R_{i,j} \log \frac{R_{i,j}}{S_{i,j}} \quad (13)$$

This directly corresponds to our *diversity* quality dimension (II), as it heavily penalizes when synthetic data misses modes present in the real distribution.

- **KL($S \parallel R$):** Kullback-Leibler divergence of the synthetic (S) distribution from the real (R) distribution measures how well the synthetic distribution stays within regions supported by real data:

$$KL(S \parallel R) = \sum_{i,j} S_{i,j} \log \frac{S_{i,j}}{R_{i,j}} \quad (14)$$

This reflects our *fidelity* quality dimension (I), penalizing synthetic data that generates unrealistic trajectories in regions with little or no real data.

²To the best of the authors ability to interpret the paper since the evaluation code is not available for DiffTraj.

1134 • **KL_{sym}**: The symmetric KL divergence balances both directional measures:
 1135

1136
$$\text{KL}_{\text{sym}} = \frac{1}{2}(\text{KL}(S \parallel R) + \text{KL}(R \parallel S)) \quad (15)$$

 1137

1138 • **JS**: Jensen-Shannon divergence provides a bounded symmetric measure:
 1139

1140
$$\text{JS}(S, R) = \frac{1}{2}\text{KL}(S \parallel M) + \frac{1}{2}\text{KL}(R \parallel M) \quad (16)$$

 1141

1142 where $M = \frac{1}{2}(S + R)$. This avoids numerical instabilities with zero probabilities and is
 1143 bounded between 0 and 1.
 1144

1145 Together, these measures address the *proportionality* quality dimension (III) - low values in both KL
 1146 directions indicate the synthetic distribution not only covers the same support as the real distribution
 1147 but also assigns similar probability mass across that support. The symmetric measures (KL_{sym} and
 1148 JS) provide balanced assessments that account for both support coverage and proportionality, with JS
 1149 offering better numerical stability when regions with zero probability are present in either distribution.
 1150

1151 **E.3 THEORETICAL JUSTIFICATION AND IMPLEMENTATION**

1152 Our implementation of KL-based measures uses a discretizations with a 256×256 grid over the
 1153 entire city. This is $\times 256$ higher resolution than what is used in DiffTraj, to enabling more precise
 1154 evaluation of fine-grained spatial patterns.
 1155

1156 KL divergence measures are theoretically well-suited for trajectory evaluation because:
 1157

1158 • They directly quantify the information loss when approximating one distribution with
 1159 another
 1160 • They are sensitive to both the support coverage and the proportionality of distributions
 1161 • The two directional variants ($\text{KL}(S \parallel R)$ and $\text{KL}(R \parallel S)$) provide complementary insights into
 1162 different failure modes
 1163

1164 **E.4 RELATION TO SAMPLE QUALITY**

1165 The combined set of measures provides a comprehensive evaluation framework addressing all five
 1166 quality dimensions from Section 4:
 1167

1168 • **Fidelity**: Length Error, Trip Error, and $\text{KL}(R \parallel S)$ all measure aspects of fidelity by ensuring
 1169 trajectories have realistic properties.
 1170 • **Diversity**: $\text{KL}(S \parallel R)$ and Pattern Score evaluate how well synthetic trajectories cover the
 1171 support of real data.
 1172 • **Proportionality**: All KL-based measures, but especially JS, capture proportionality by
 1173 measuring distributional similarity.
 1174 • **Usefulness**: TSTR directly measures usefulness for a downstream prediction task.
 1175 • **Generalization**: See Section E.6.
 1176

1177 **E.5 LIMITATIONS ANALYSIS**

1178 The DiffTraj measures provide a more coarse-grained evaluation using 16×16 spatial discretization,
 1179 while our KL-based measures use 256×256 grids for finer-grained assessment. The $\times 256$ higher
 1180 resolution allows more precise evaluation of the model's ability to capture detailed spatial distributions,
 1181 especially important in urban environments where road networks create complex movement patterns.
 1182

1183 By combining these measures with TSTR, our benchmark provides a comprehensive evaluation that
 1184 addresses both distribution-level properties and practical utility of the generated trajectories, spanning
 1185 all important dimensions of synthetic data quality discussed in Section 4.
 1186

1188
1189

E.6 EVALUATING GENERALIZATION

1190
1191

Our evaluation framework extends beyond traditional in-distribution testing to assess different types of generalization capabilities:

1192

- *In-distribution generalization*: Unlike predictive models where test sets evaluate prediction accuracy, unconditional generative models require evaluating whether they capture the full distribution rather than memorizing training examples. For trajectory generation, we directly compare distribution-level properties between synthetic and real data instead of using held-out test sets. Using a held-out test would require massive data amounts and/or an inverted data split, e.g. 20% train, 80% test to sufficiently assess support coverage and proportionality. This is currently infeasible in practice. Also, current models still struggle to even replicate (memorize) the training data in unconditional generation, making this less of a relevant problem in the field for now.
- *Out-of-distribution generalization*: A distinctive feature of our benchmark is its assessment of spatial transfer capability, evaluating a model’s ability to generate plausible trajectories for geographic regions not represented in training data:
 - *Intra-city generalization* (Section 4.3): Tests whether models can generate trajectories for unseen areas within the same city
 - *City-to-city generalization* (Section 4.3): Tests whether models can transfer knowledge between entirely different cities

1209

The evaluation measures take on different significance when assessing generalization. $KL(S\|R)$ becomes particularly important for out-of-distribution evaluation as it tests whether the synthetic distribution covers the full support of the real distribution in previously unseen regions. Meanwhile, $KL(R\|S)$ reveals whether the model avoids generating implausible trajectories in new environments. The Pattern Score helps determine if important spatial areas in the new environment are captured, while TSTR directly measures the usefulness of the generated data for downstream tasks in these new settings.

1216

Our experiments in Sections 4.3 demonstrate how this comprehensive evaluation identifies genuine generalization capabilities rather than mere memorization. The combination of measures reveals whether a model has learned transferable spatial dynamics or is simply reproducing patterns from its training data. This evaluation approach provides a more rigorous assessment of generalization than traditional machine learning benchmarks, which typically only test in-distribution generalization.

1221

1222

F ADDITIONAL RESULTS

1224

In Section 4 we present the results. Here we first provide additional visualization, focusing on more detailed view of and also providing a closer-look at the synthetic data generated. We also provide full tables as a complement to the tables provide figures that are averaged across datasets.

1228

1229

F.1 ADDITIONAL VISUALIZATION

1230

The visualizations presented in this section supports the quantitative results in Section 4.1 of the main paper. Specifically, Figures 6–8 demonstrate TDDM’s ability to generate synthetic trajectories that closely match the spatial distribution patterns of the original data across all three datasets (Geolife, Porto, and Cabspotting).

1234

These visualizations highlight two key aspects of TDDM’s performance claimed in the main paper: (1) the high fidelity of individual generated trajectories, which follow road networks and maintain realistic movement patterns; and (2) the proportionality of the generated distribution, as shown by the similarity between the heatmaps of real and synthetic data. The comparison with baseline methods in Figures 2, 4 and 5 visually confirms TDDM’s strong performance reported in Table 7.

1240

1241

The generalization capabilities of TDDM are further illustrated in Figures 12–14, which show how our model trained on a limited portion of each city can generate high-quality trajectories for previously unseen areas. Additionally, Figures 9–11 provide a closer look at trajectories in several regions

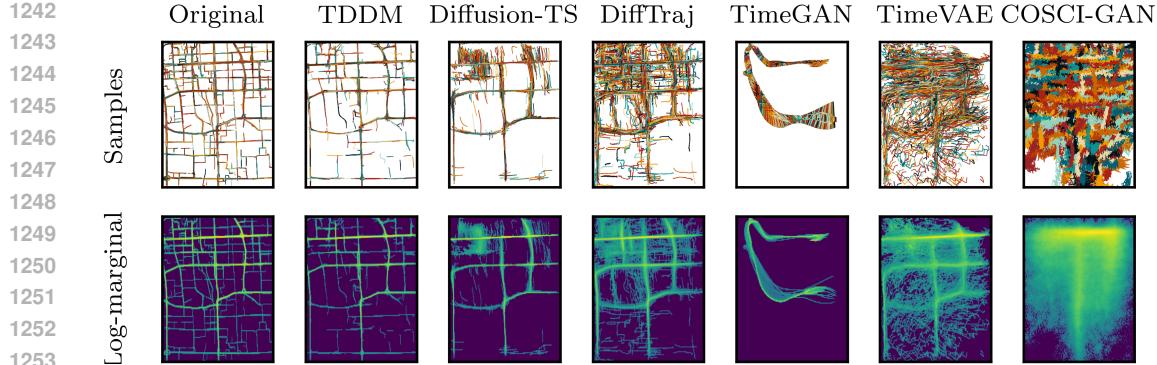


Figure 4: Comparison between original and synthetic trajectories for Geolife dataset. First row shows individual trajectory samples while bottom shows log-density heatmaps of all observations. The synthetic data of the proposed model, TDDM, (second column) most closely matches both the individual trajectory patterns and overall density distribution of the original data.

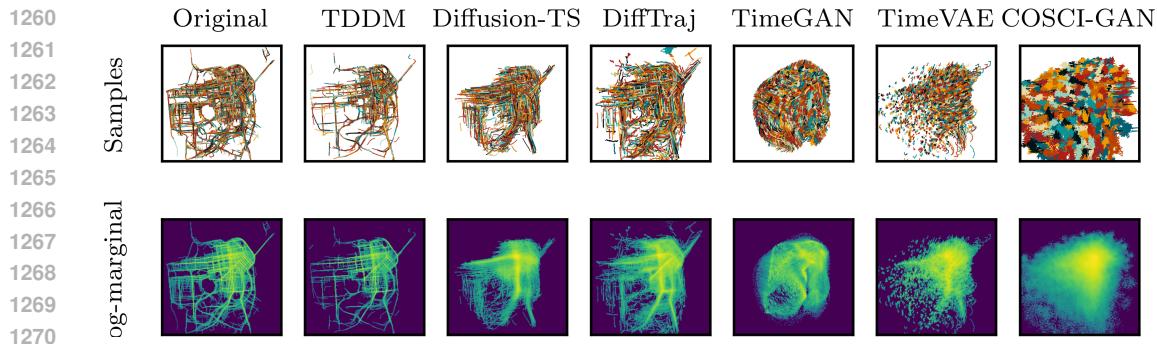


Figure 5: Comparison between original and synthetic trajectories for Cabspotting dataset. First row shows individual trajectory samples while bottom shows log-density heatmaps of all observations. The synthetic data of the proposed model, TDDM, (second column) most closely matches both the individual trajectory patterns and overall density distribution of the original data.

1278 chosen at random, across the three datasets, demonstrating the consistent quality of our generated
 1279 data compared to real trajectories and baseline methods.

1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320

Original samples



1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342

Synthetic samples



1343
1344
1345
1346
1347
1348
1349

Original log-marginal



Synthetic log-marginal

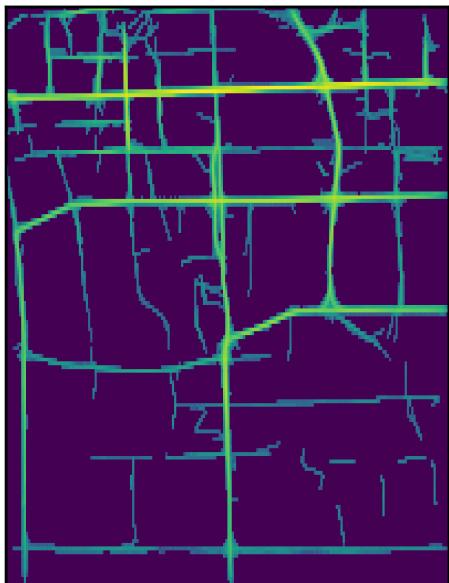


Figure 6: Detailed visualization for Geolife dataset. Comparison between training data (top) and synthetic trajectories from our method (bottom). Left panels show individual trajectory samples while right panels show log-density heatmaps of all points in the dataset. The synthetic samples closely match both the individual trajectory patterns and overall density distribution of the original data.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

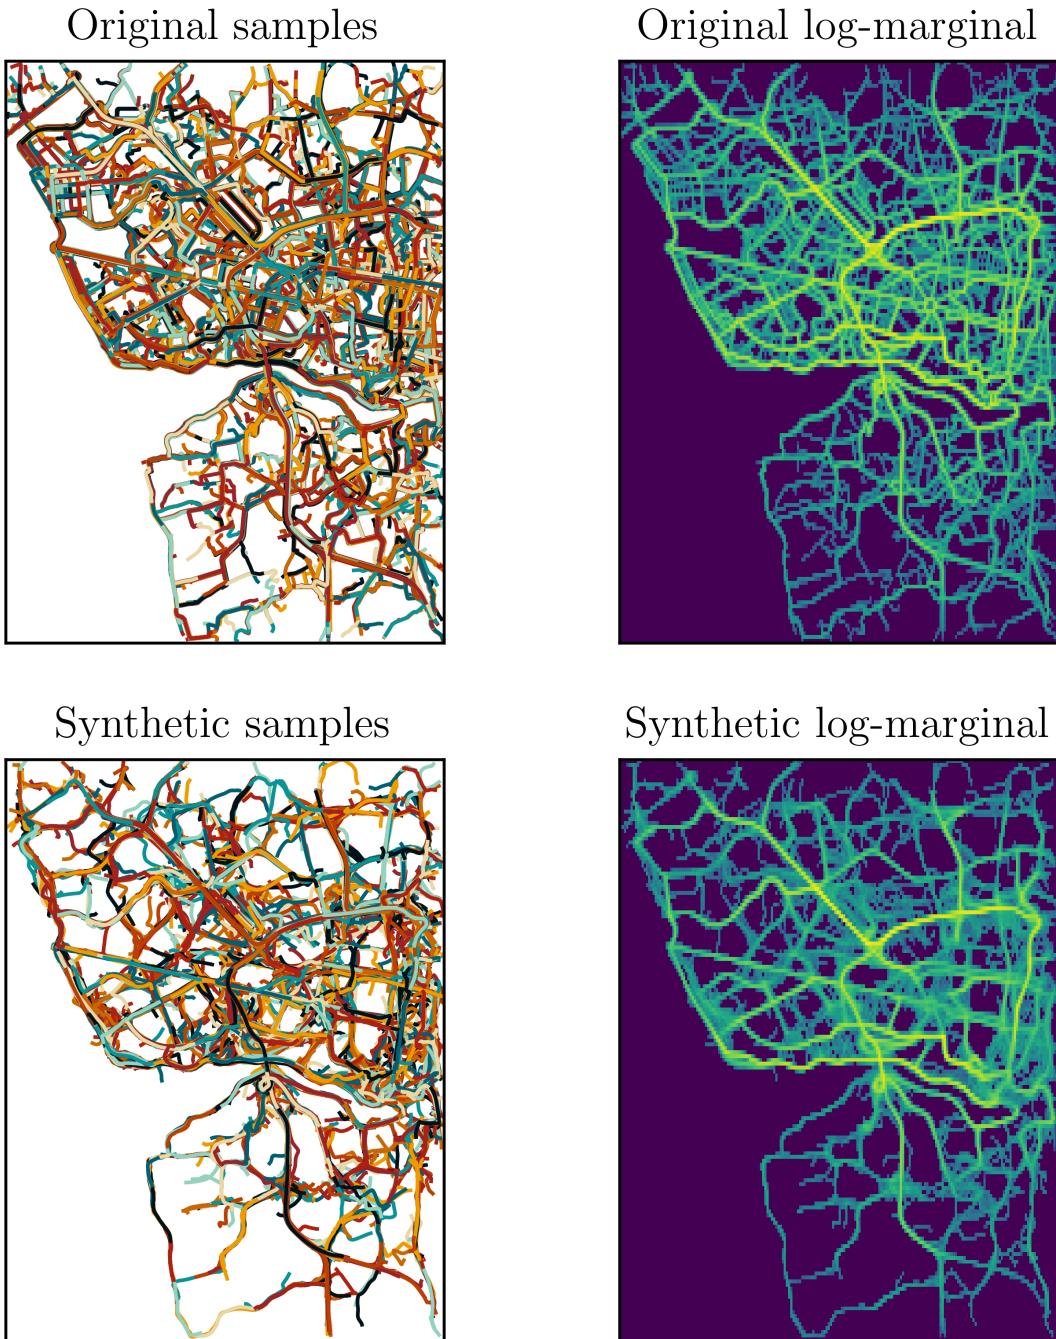
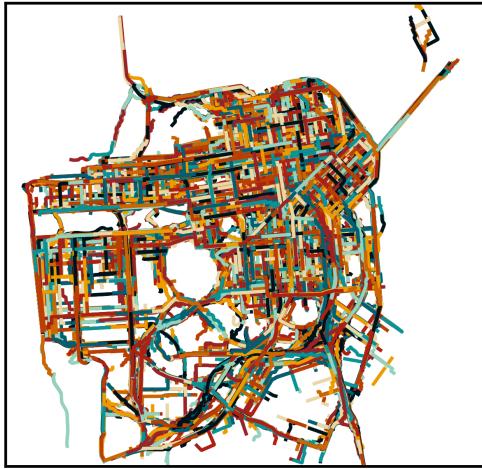


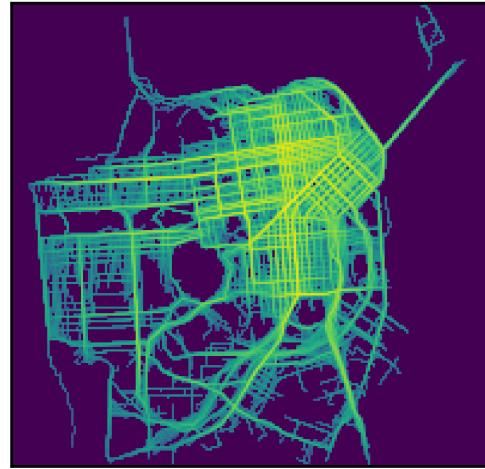
Figure 7: Detailed visualization for Porto dataset. Comparison between training data (top) and synthetic trajectories from our method (bottom). Left panels show individual trajectory samples while right panels show log-density heatmaps of all points in the dataset. The synthetic samples closely match both the individual trajectory patterns and overall density distribution of the original data.

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426

Original samples

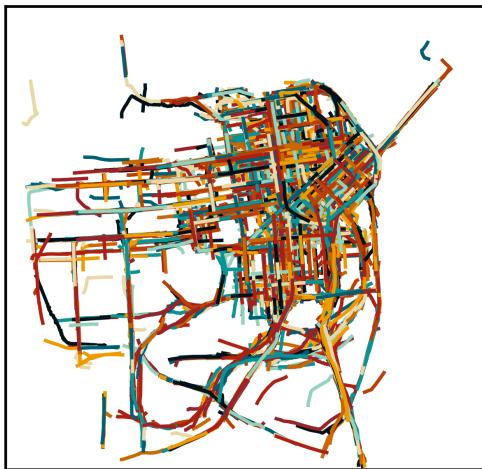


Original log-marginal

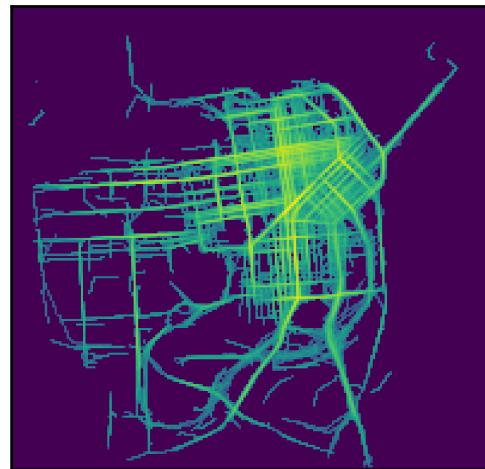


1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448

Synthetic samples



Synthetic log-marginal



1449
1450
1451
1452
1453
1454
1455
1456
1457

Figure 8: Detailed visualization for Cabspotting dataset. Comparison between training data (top) and synthetic trajectories from our method (bottom). Left panels show individual trajectory samples while right panels show log-density heatmaps of all points in the dataset. The synthetic samples closely match both the individual trajectory patterns and overall density distribution of the original data.

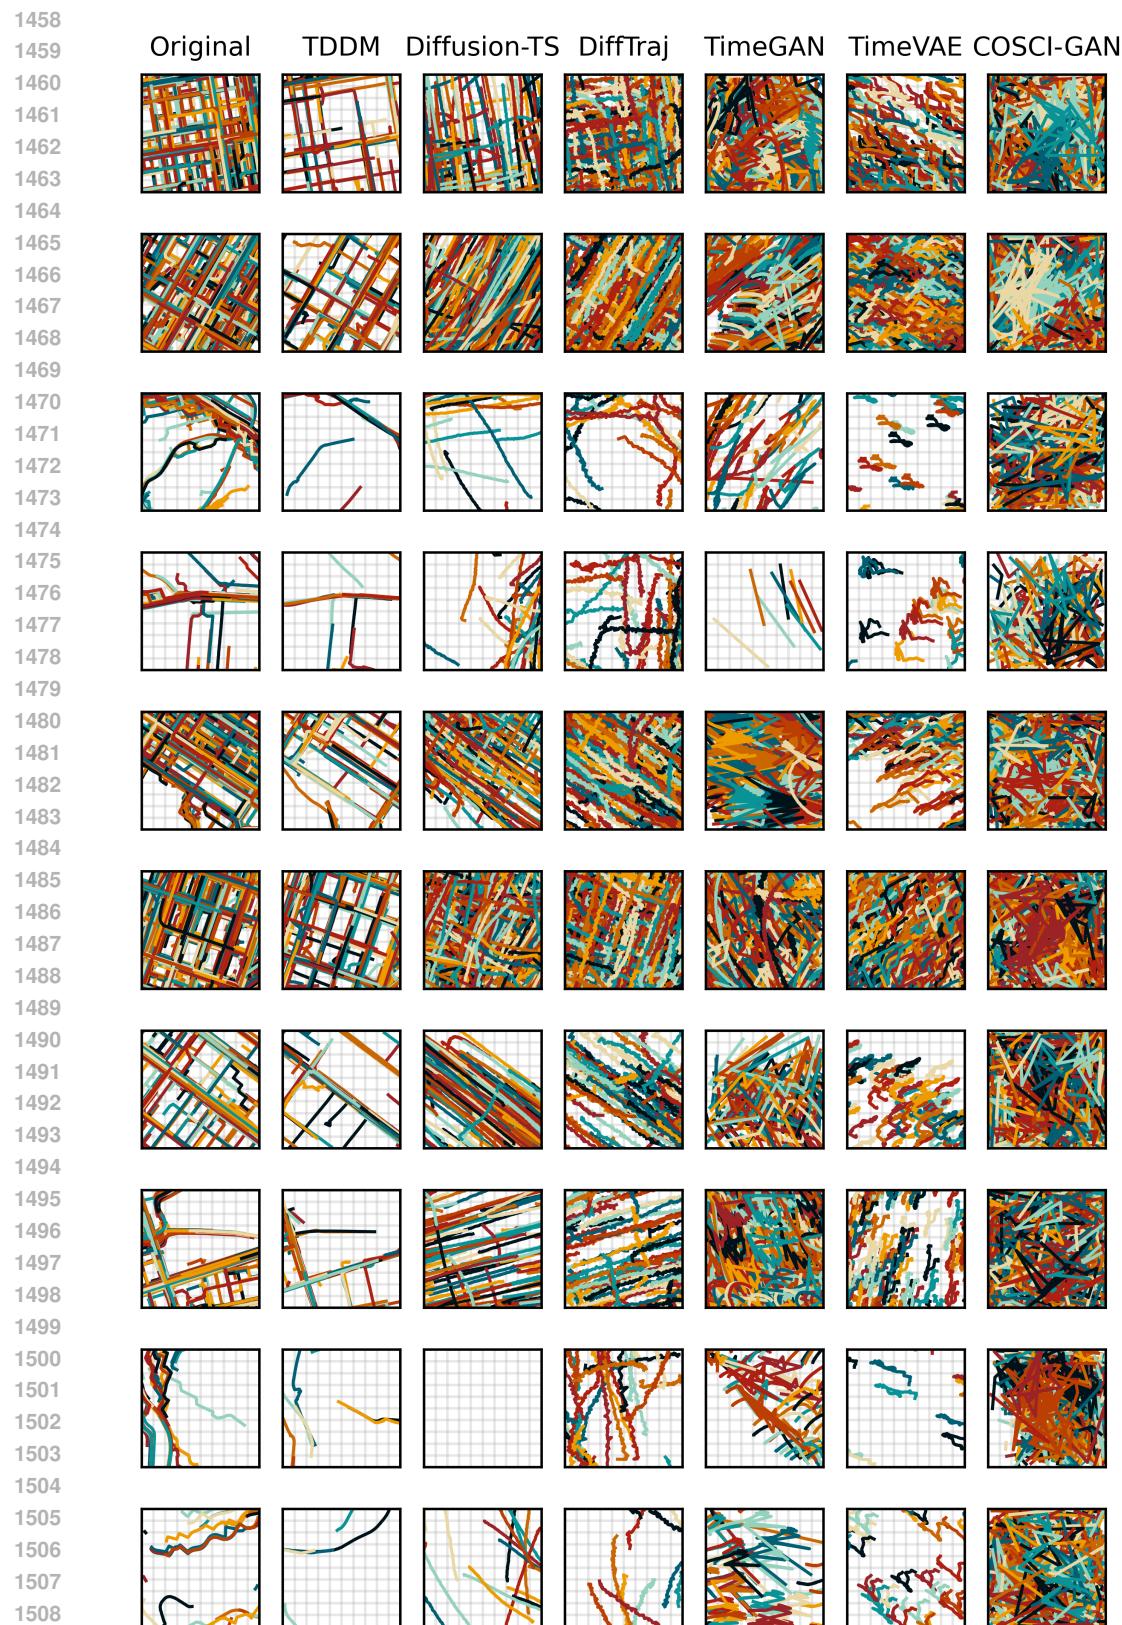


Figure 9: Random samples from training data and synthetic data across 11 different regions chosen at random, all from Cabspotting

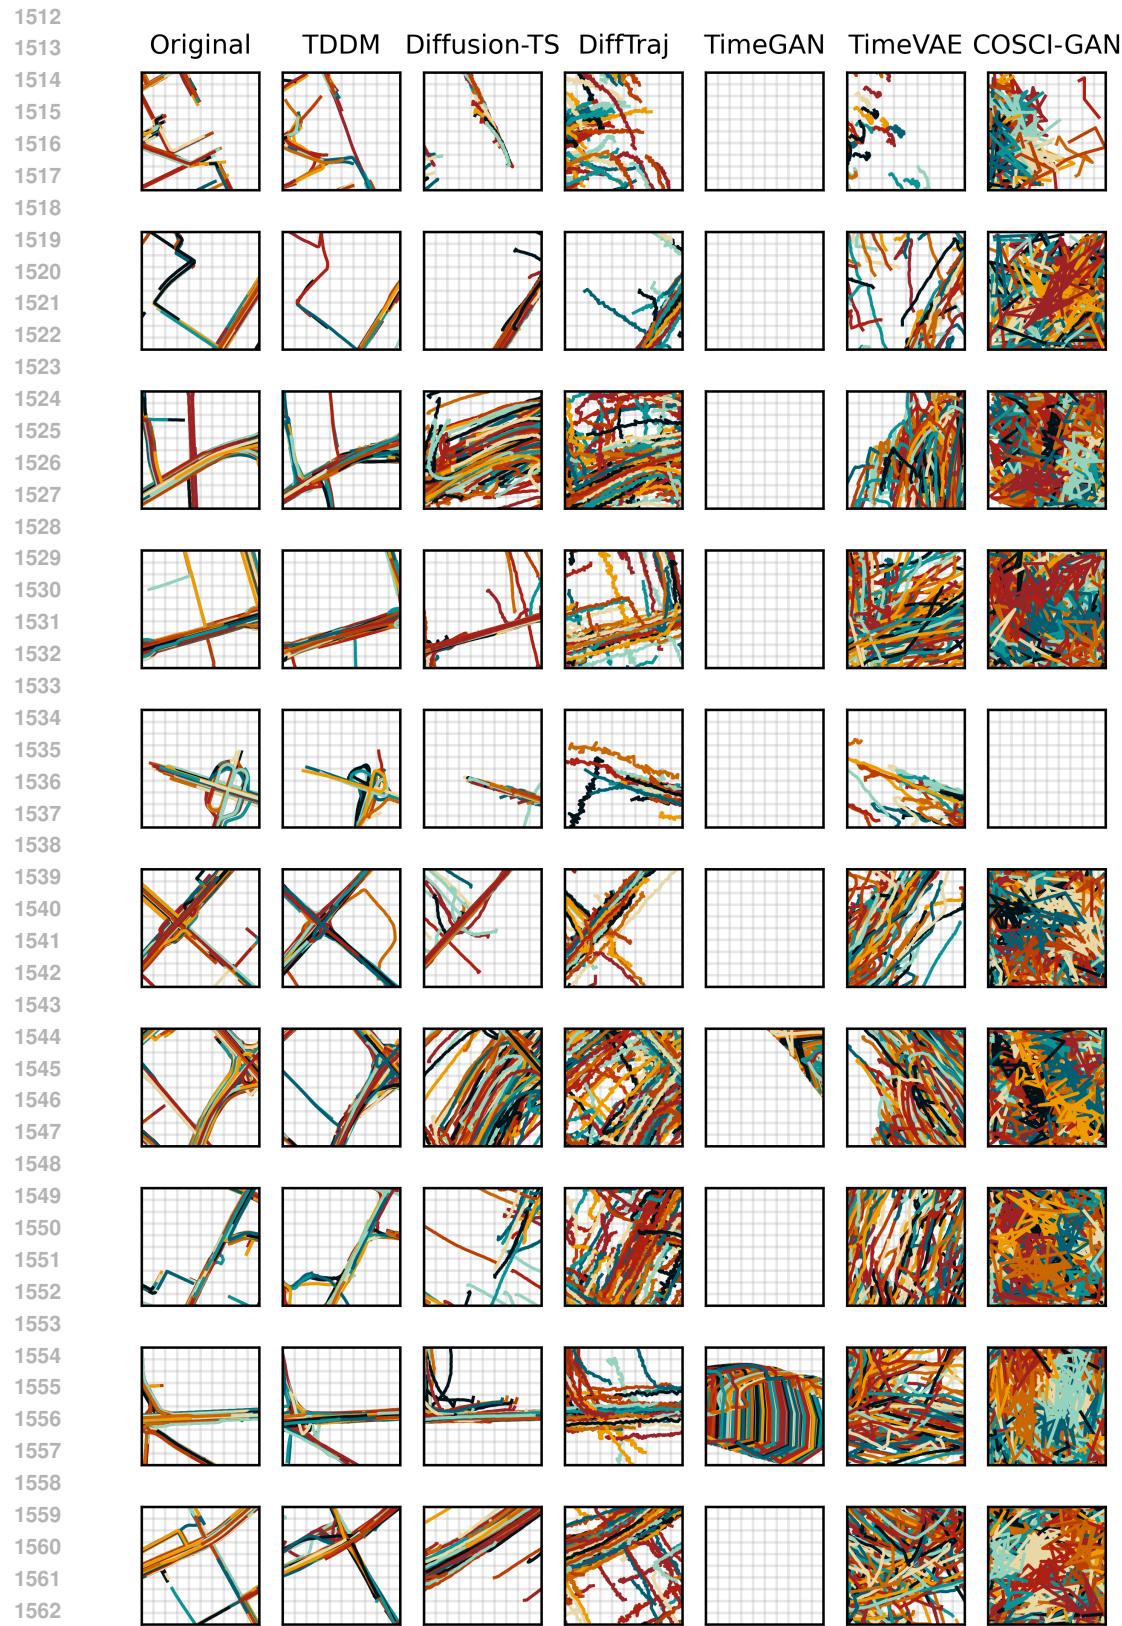


Figure 10: Random samples from training data and synthetic data across 11 different regions chosen at random, all from Geolife

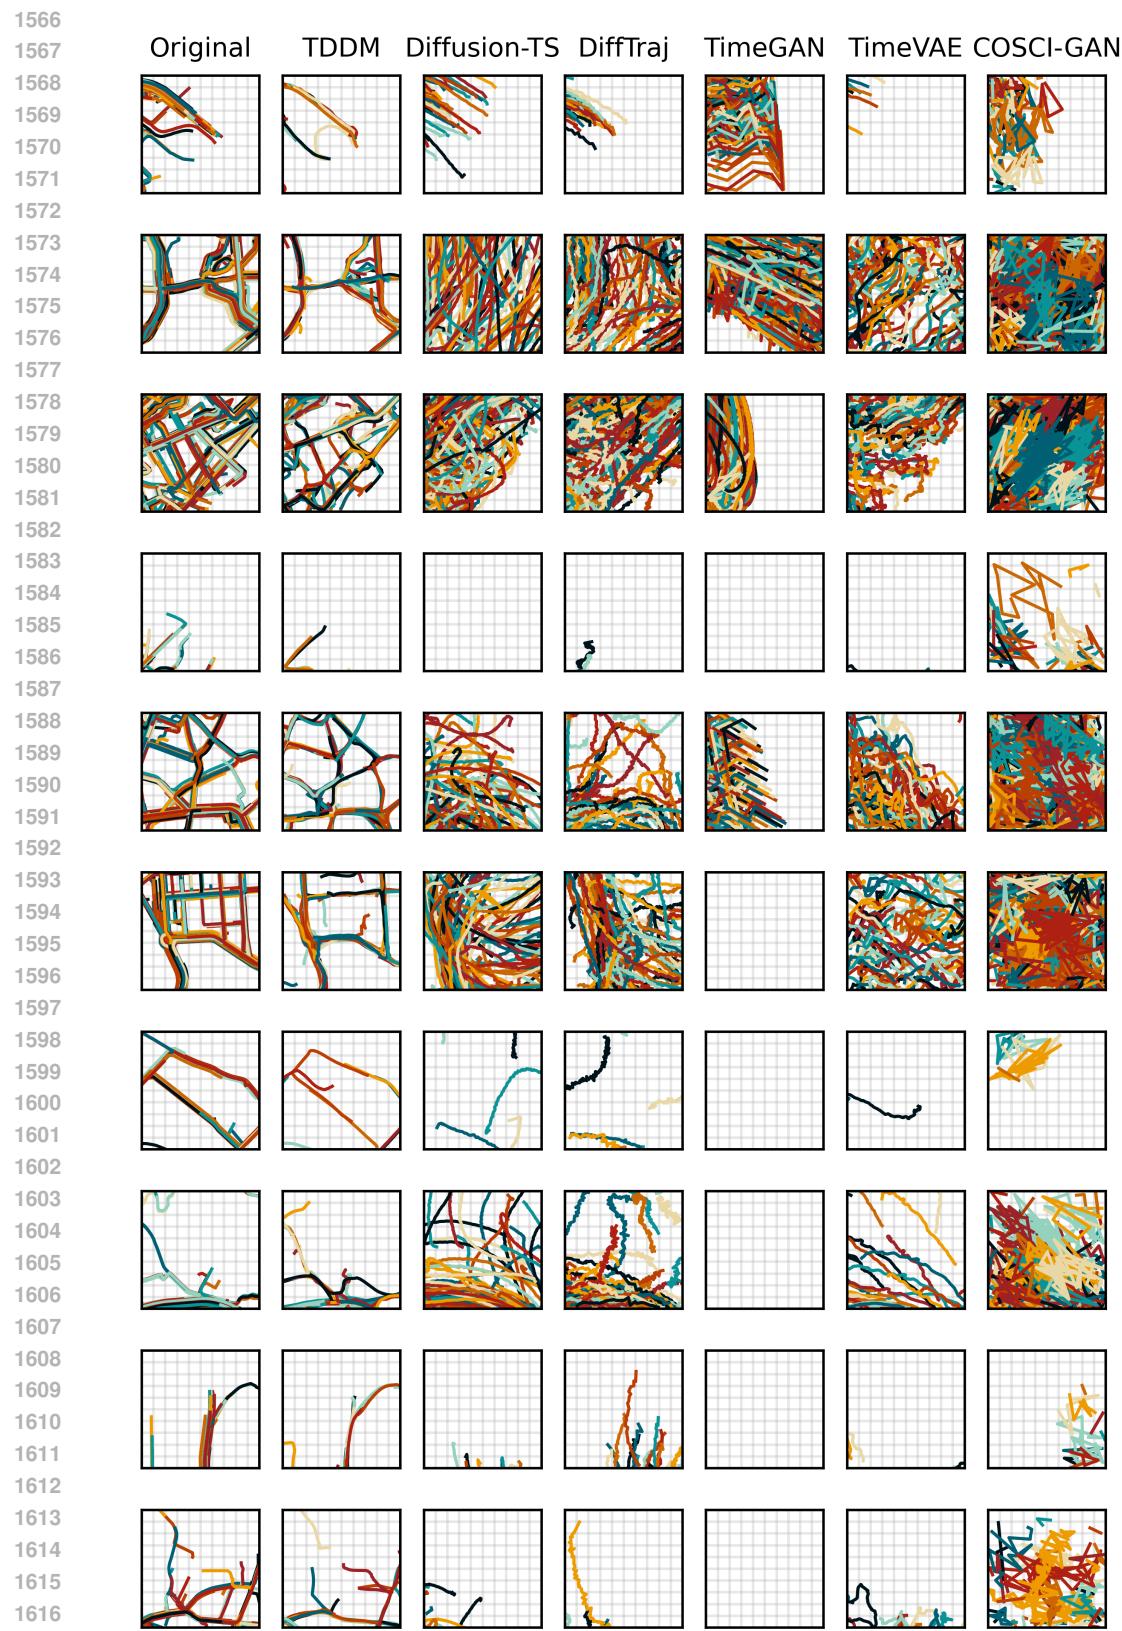


Figure 11: Random samples from training data and synthetic data across 10 different regions chosen at random, all from Porto.

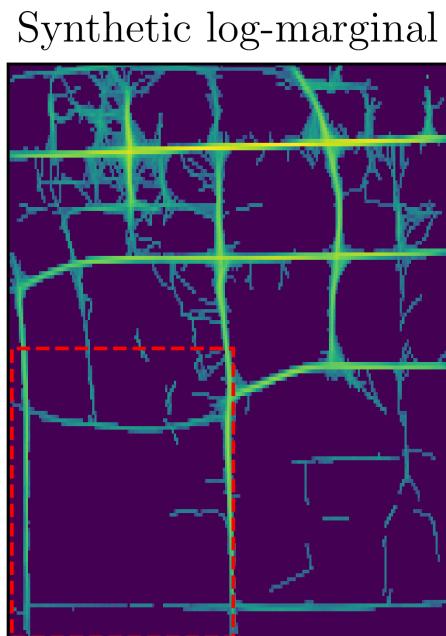
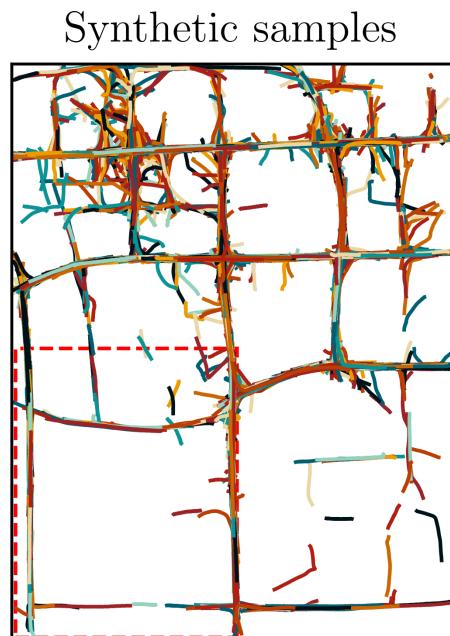
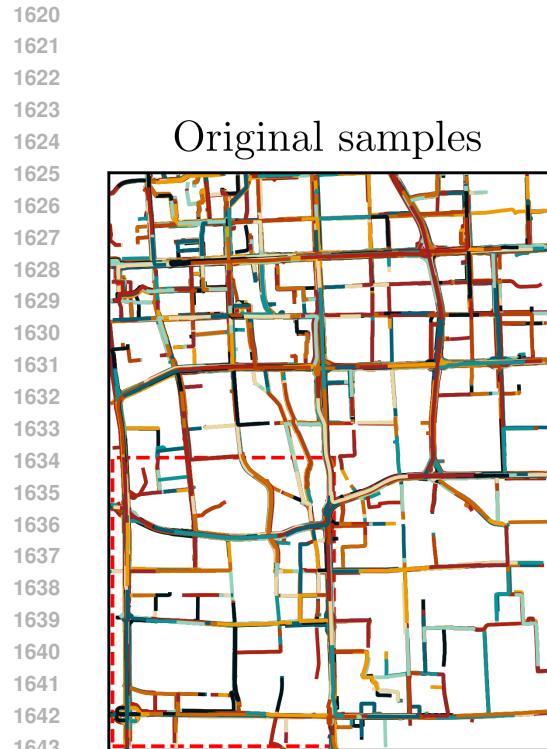
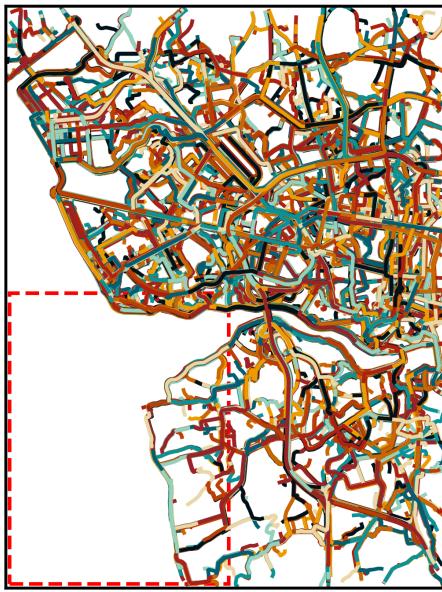
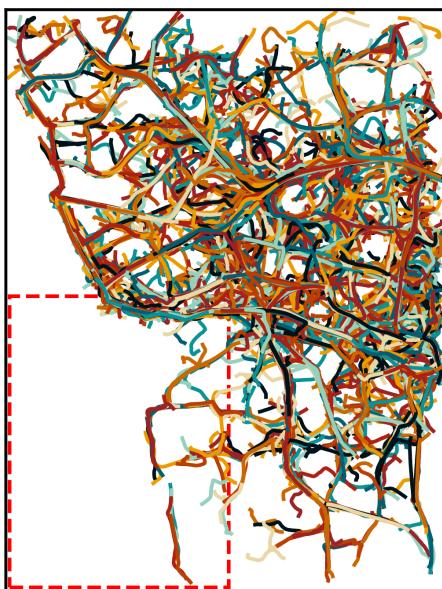


Figure 12: Generalization experiment. The model is trained on the lower left quadrant and used to generate data on the remaining geographical area. *Top left*: Data from Geolife, the lower-left quadrant of which used for training. *Top right*: heatmap of training data, *Bottom left*: synthetic trajectories. *Bottom right*: heatmap of the synthetic data.

1674
 1675
 1676
 1677
 1678 Original samples
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697

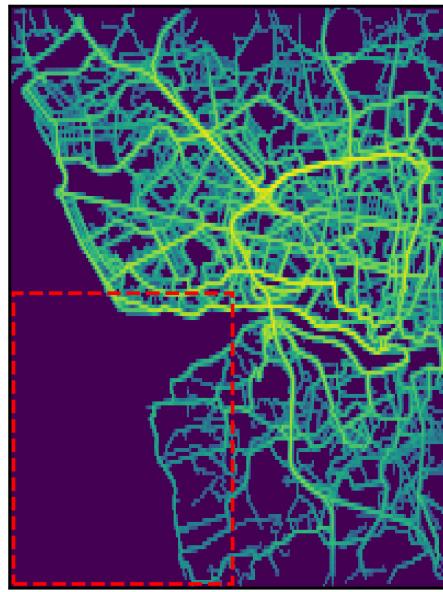


1698
 1699
 1700 Synthetic samples
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719

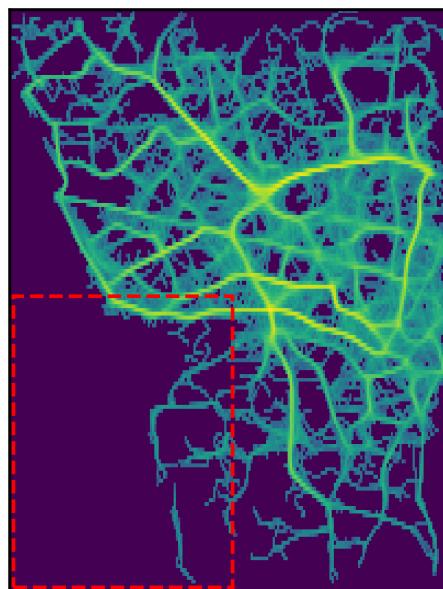


1720
 1721
 1722
 1723
 1724 Figure 13: Generalization experiment. The model is trained on the top-left quadrant and used to
 1725 generate data on the remaining geographical area. *Top left*: Data from Porto, the top-left quadrant
 1726 of which used for training. *Top right*: heatmap of training data, *Bottom left*: synthetic trajectories.
 1727 *Bottom right*: heatmap of the synthetic data.

Original log-marginal



Synthetic log-marginal



1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781

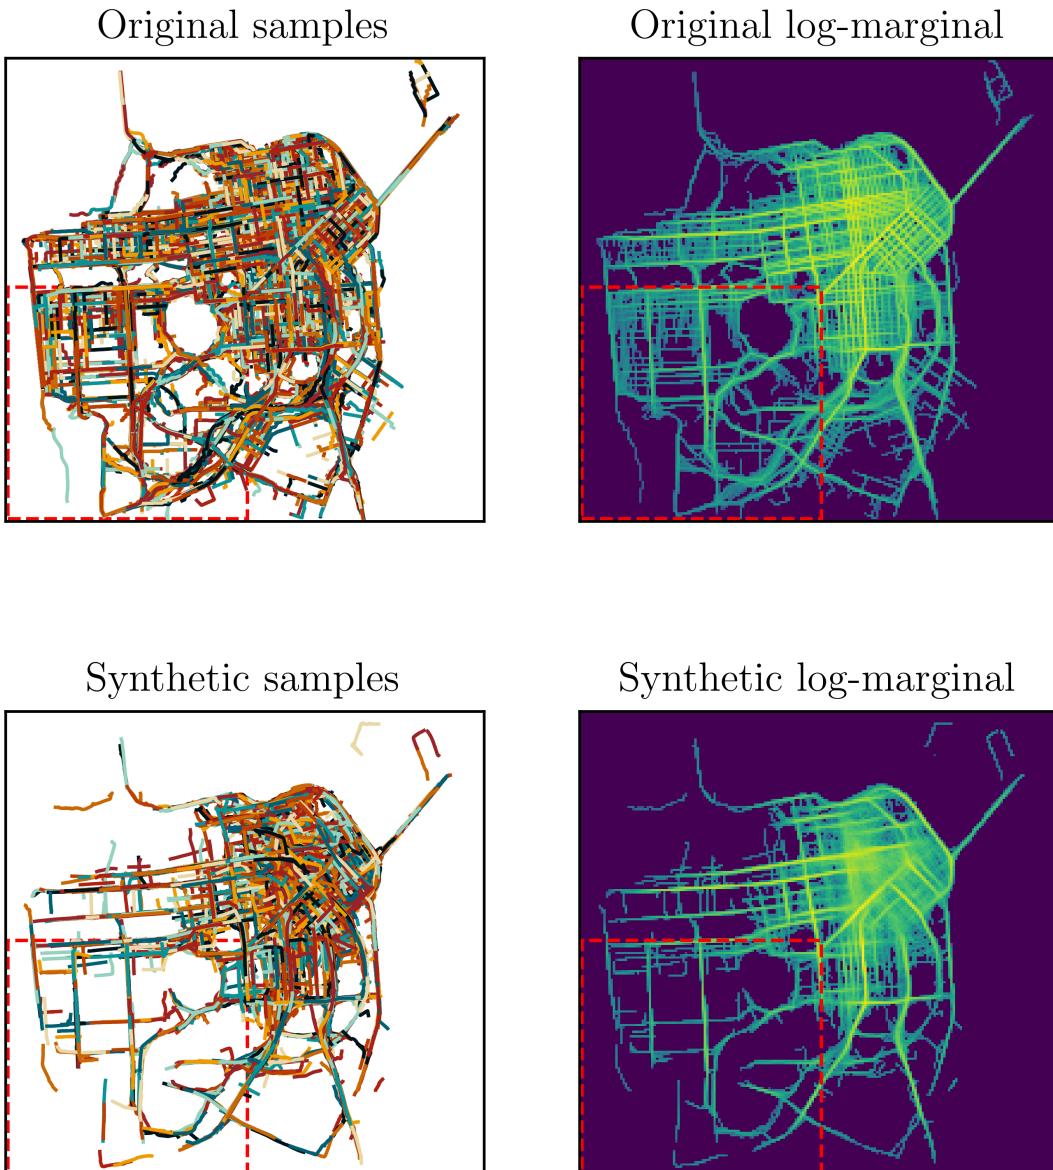
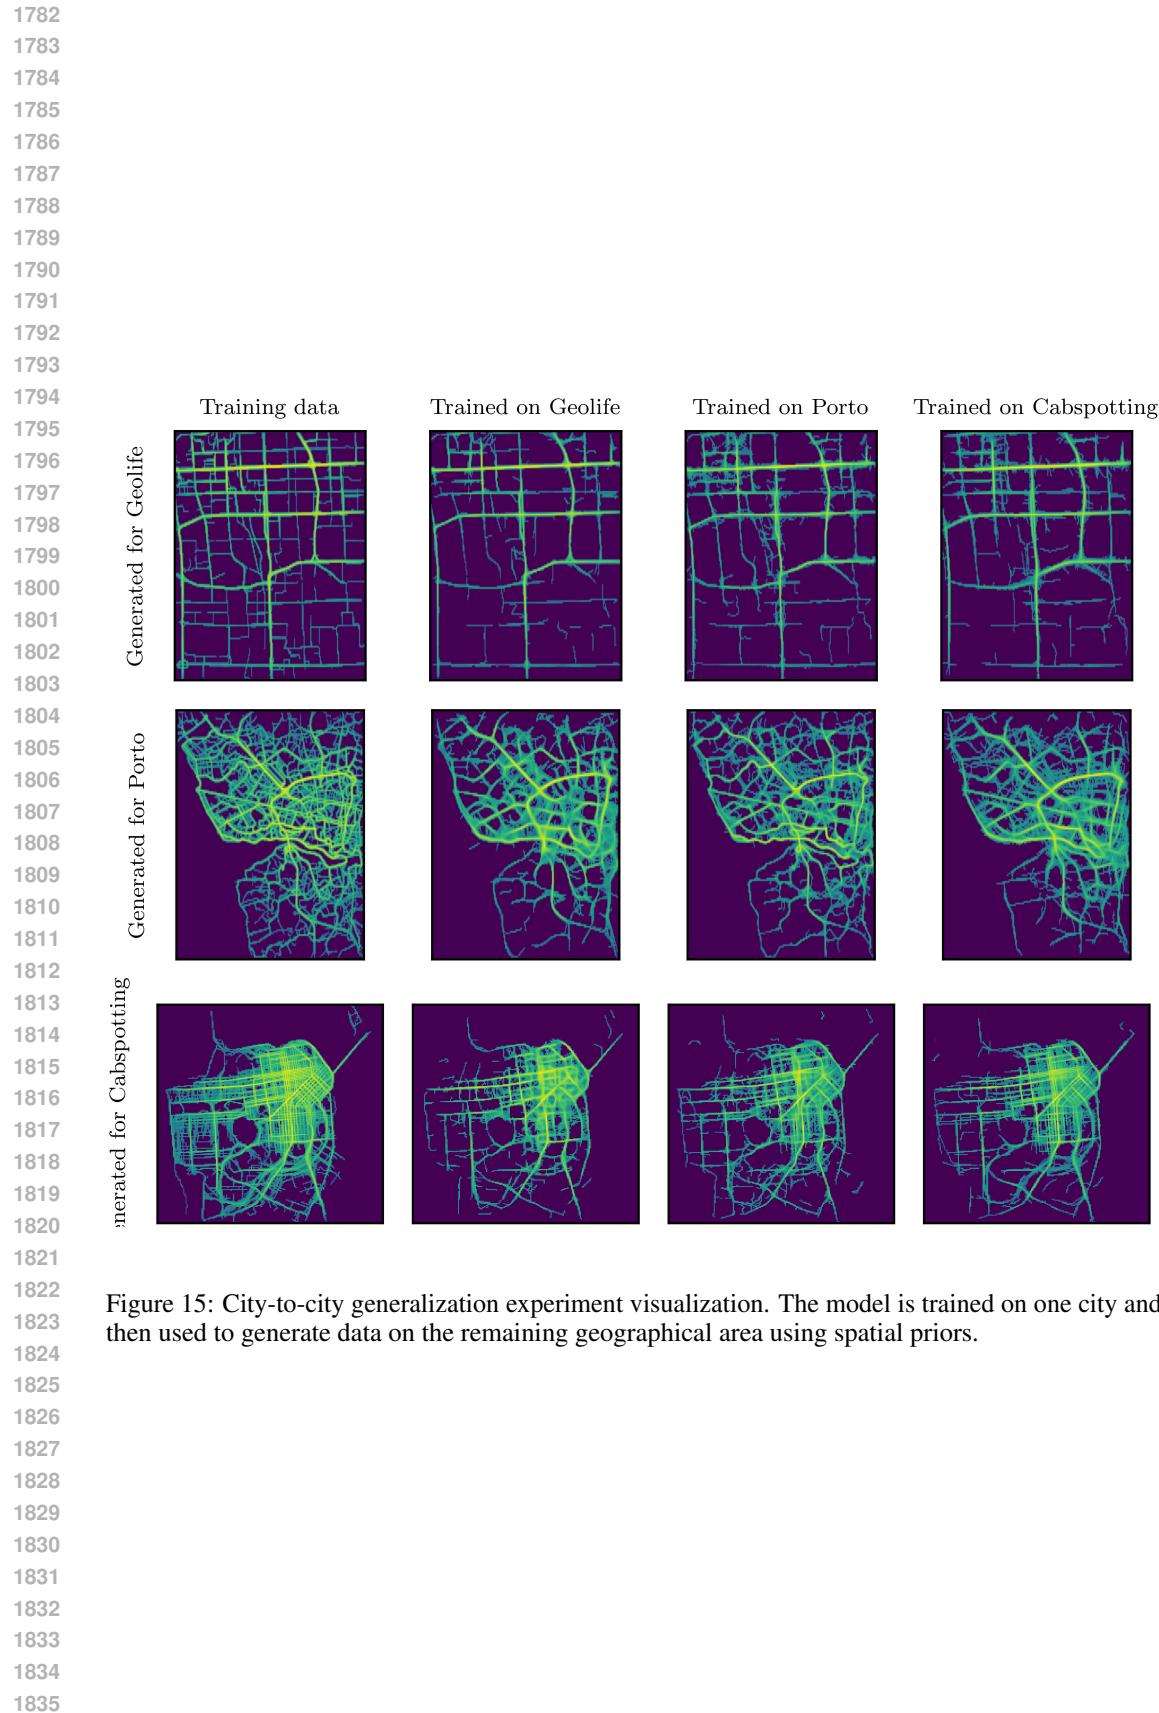


Figure 14: Generalization experiment. The model is trained on the lower left quadrant and used to generate data on the remaining geographical area. *Top left*: Data from Cabspotting, the lower-left quadrant of which used for training. *Top right*: heatmap of training data, *Bottom left*: synthetic trajectories. *Bottom right*: heatmap of the synthetic data.



1836
1837

F.2 LENGTH DISTRIBUTION ANALYSIS

1838
1839
1840
1841
1842

To provide deeper insight into the Length Error measure described in Section 4, we visualize the length distributions (distance between consecutive trajectory points) for both training and synthetic data across all experimental settings. Each histogram shows the normalized distribution of Euclidean distances between consecutive points, with the Jensen-Shannon divergence between these distributions reported in the top-right corner.

1843
1844
1845
1846

These visualizations demonstrate TDDM’s ability to accurately match the fine-grained movement patterns of real trajectories. The close alignment between training (teal) and synthetic (red) distributions across diverse experimental conditions validates the quantitative Length Error results presented in the main paper.

1847

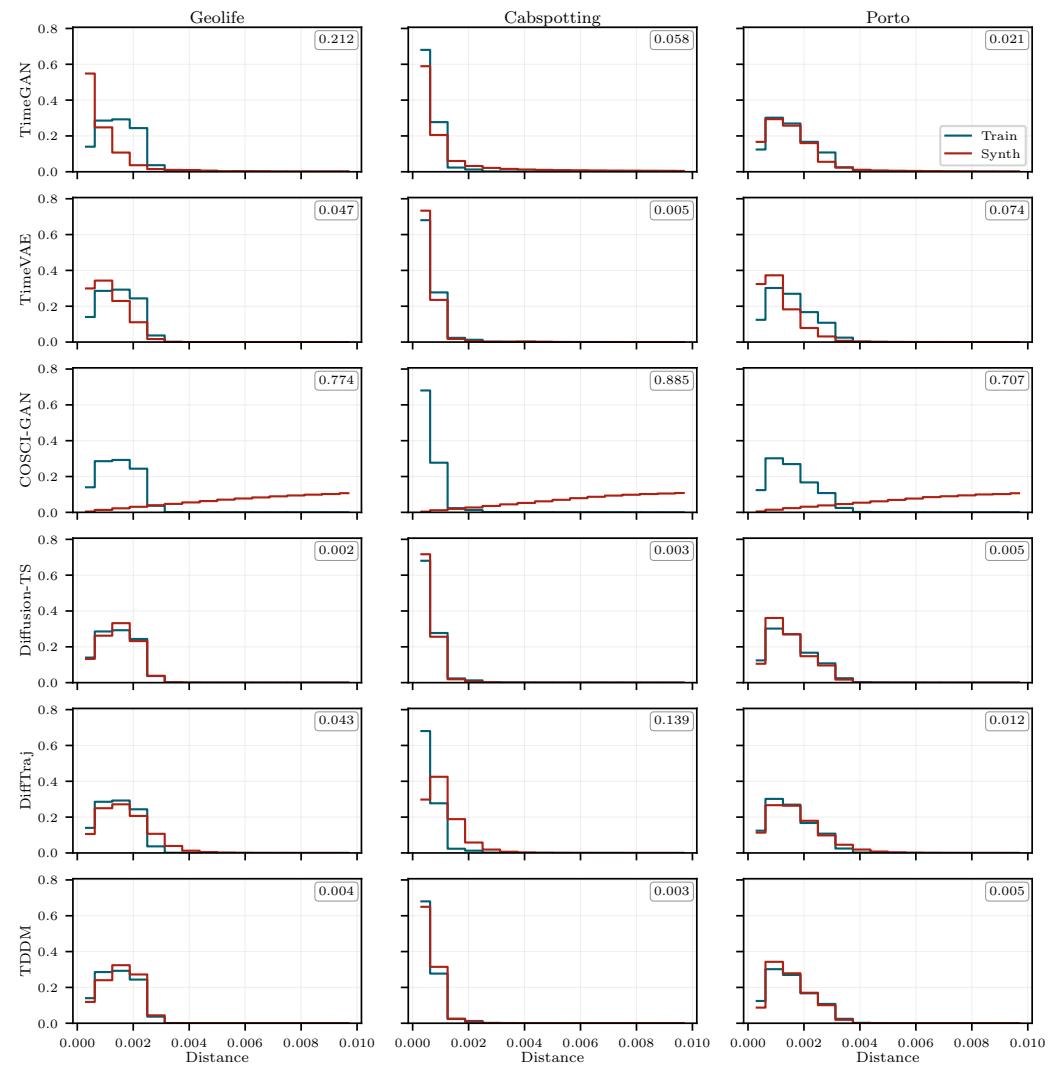
1883
1884
1885
1886
1887
1888
1889

Figure 16: Length distribution comparison for unconditional generation benchmark. Each subplot shows the distribution of distances between consecutive trajectory points for training data (teal) and synthetic data (red) across three datasets (Geolife, Cabspotting, Porto) and six methods. The Jensen-Shannon divergence is shown in the top-right corner of each subplot.

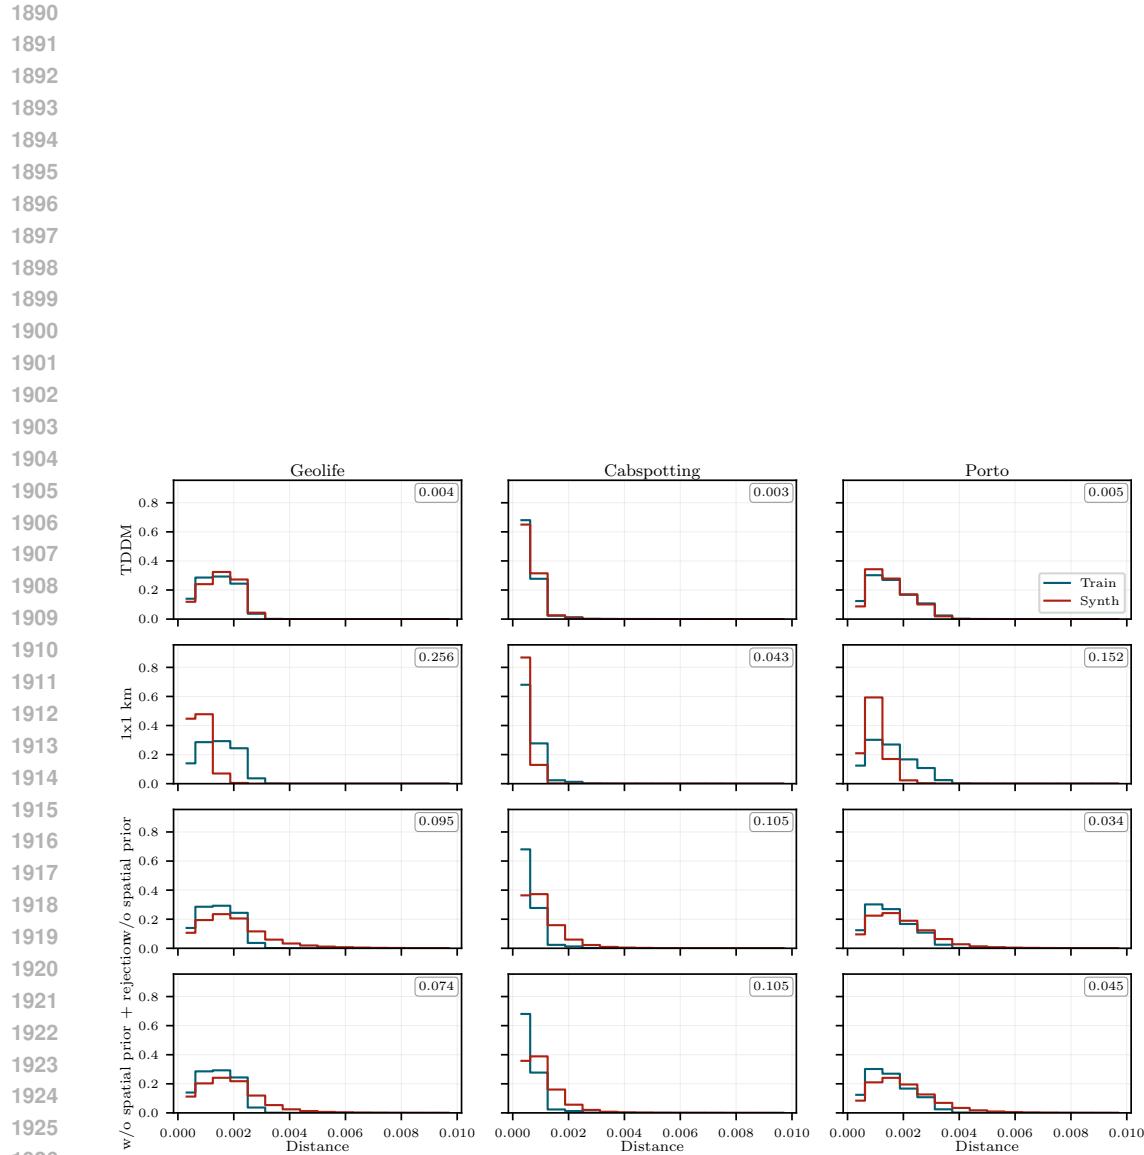


Figure 17: Length distribution comparison for ablation study. Shows the impact of different architectural choices on TDDM’s ability to match the length distribution of training data.

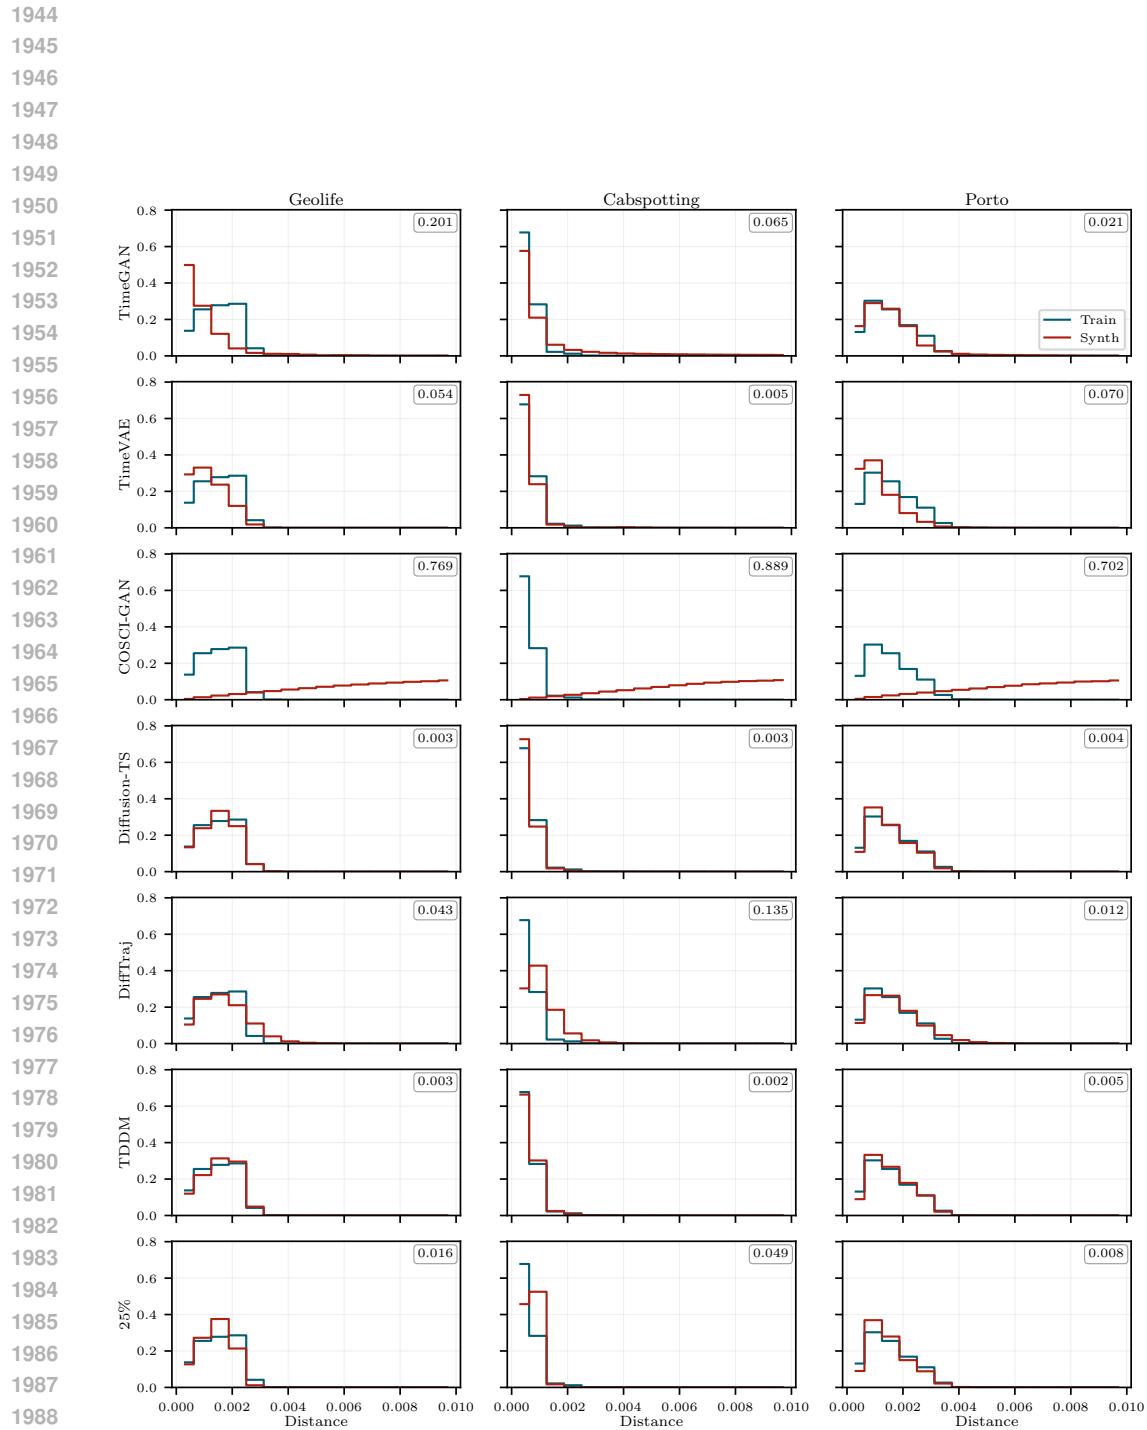


Figure 18: Length distribution comparison for intra-city generalization experiments. Models are trained on 25% of the city and evaluated on the full area.

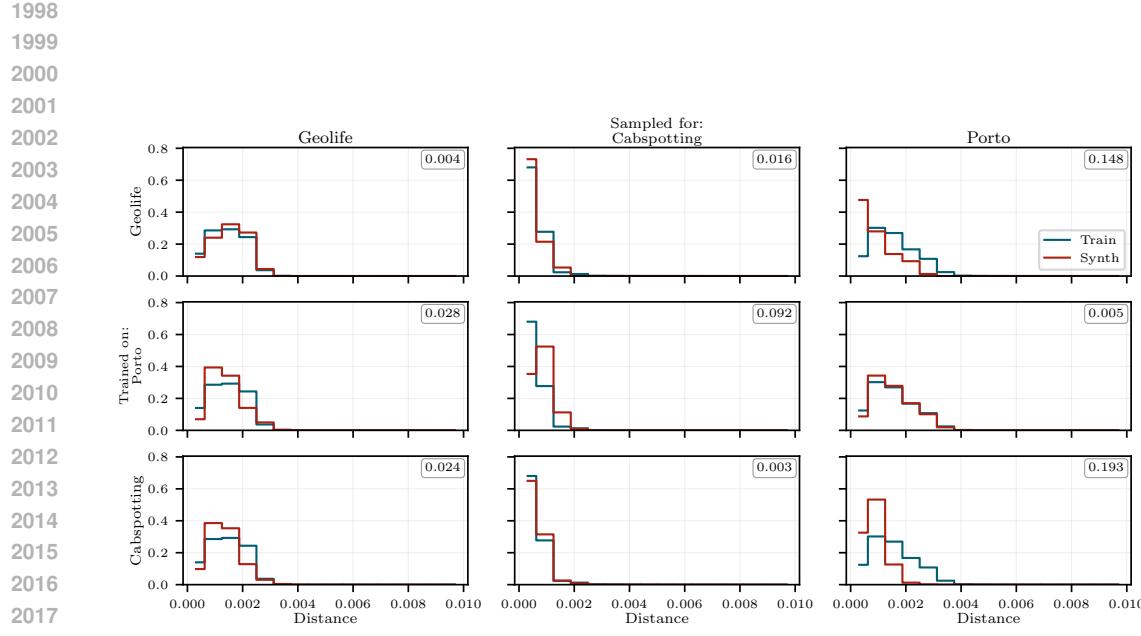


Figure 19: Length distribution comparison for city-to-city generalization experiments. Each row shows a model trained on one city and evaluated on all three cities.

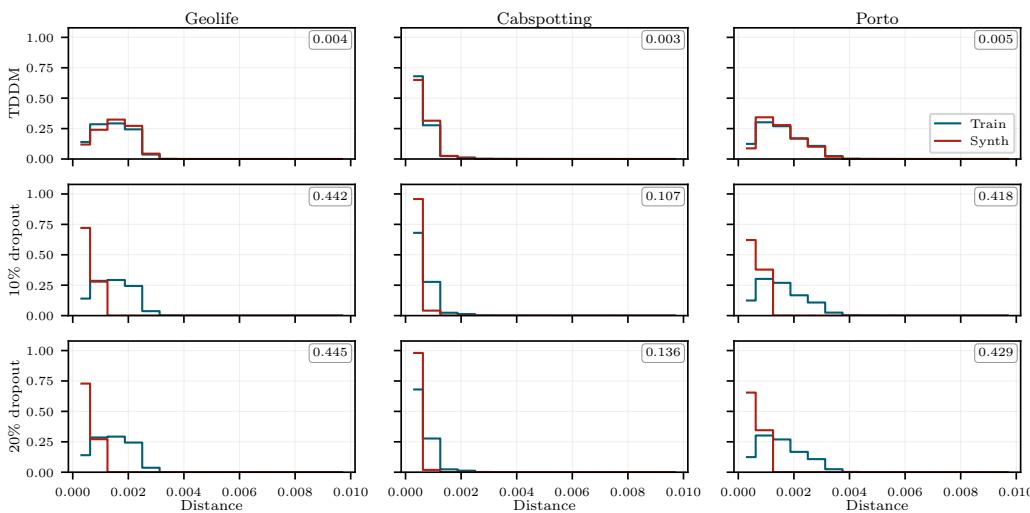


Figure 20: Length distribution comparison for map dropout robustness experiments. Shows TDDM's performance with varying levels of missing road network information.

2106 F.3 QUANTITATIVE RESULTS
21072108 Complementing the summary tables, we list the tables with the quantitative results.
21092110 More specifically:
21112112

- *Large-scale unconditional trajectory generation* (4.1), Table 7.
- *Ablation study* (4.2), Table 8.
- *Intra-city* (4.3), Table 11.
- *City-to-city* (4.3), Table 12.
- *Map matching ablation*, Table 9.
- *Robustness to map dropout*, Table 10.

2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159

2160

2161

2162

2163

2164

2165 Table 7: Evaluation of different models’ performance across several datasets and measures.

Measure	Model	Geolife	Porto	Cabspotting
TSTR (↓)	TimeGAN	0.052 ± 0.031	0.038 ± 0.025	0.020 ± 0.012
	TimeVAE	0.017 ± 0.010	0.016 ± 0.007	0.021 ± 0.011
	COSCI-GAN	0.023 ± 0.006	0.021 ± 0.005	0.024 ± 0.009
	Diffusion-TS	<u>0.013 ± 0.009</u>	<u>0.010 ± 0.005</u>	0.021 ± 0.008
	DiffTraj	<u>0.013 ± 0.004</u>	0.014 ± 0.004	0.013 ± 0.007
KL($S \parallel R$) (↓)	TDDM	0.010 ± 0.005	0.009 ± 0.004	<u>0.015 ± 0.008</u>
	TimeGAN	5.566	3.606	1.933
	TimeVAE	2.786	2.688	1.615
	COSCI-GAN	4.650	2.910	1.578
	Diffusion-TS	<u>0.730</u>	2.114	1.340
KL($R \parallel S$) (↓)	DiffTraj	1.498	<u>2.025</u>	<u>1.259</u>
	TDDM	0.202	0.437	0.265
	TimeGAN	3.823	2.632	1.302
	TimeVAE	1.236	1.445	1.124
	COSCI-GAN	2.453	1.696	1.070
KL _{sym} (↓)	Diffusion-TS	<u>0.598</u>	1.135	0.999
	DiffTraj	0.767	<u>1.023</u>	<u>0.817</u>
	TDDM	0.177	0.336	0.246
	TimeGAN	4.695	3.119	1.618
	TimeVAE	2.011	2.067	1.369
JS (↓)	COSCI-GAN	3.552	2.303	1.324
	Diffusion-TS	<u>0.664</u>	1.625	1.169
	DiffTraj	1.132	<u>1.524</u>	<u>1.038</u>
	TDDM	0.190	0.386	0.255
	TimeGAN	0.489	0.430	0.271
KL _{speed} (↓)	TimeVAE	0.294	0.327	0.241
	COSCI-GAN	0.486	0.364	0.240
	Diffusion-TS	<u>0.114</u>	0.264	0.216
	DiffTraj	0.181	<u>0.252</u>	<u>0.195</u>
	TDDM	0.040	0.079	0.056
Density (↓)	TimeGAN	0.324	0.361	0.710
	TimeVAE	0.351	0.240	0.083
	COSCI-GAN	6.581	5.216	7.592
	Diffusion-TS	0.038	0.015	<u>0.053</u>
	DiffTraj	<u>0.031</u>	0.035	0.312
Trip (↓)	TDDM	0.011	<u>0.020</u>	0.007
	TimeGAN	0.433	0.252	0.088
	TimeVAE	0.063	0.028	0.038
	COSCI-GAN	0.264	0.078	0.060
	Diffusion-TS	<u>0.035</u>	<u>0.021</u>	<u>0.031</u>
Length (↓)	DiffTraj	0.046	0.023	<u>0.031</u>
	TDDM	0.023	0.019	0.014
	TimeGAN	0.510	0.346	0.115
	TimeVAE	0.078	0.049	0.040
	COSCI-GAN	0.295	0.096	0.082
Pattern (↑)	Diffusion-TS	<u>0.049</u>	<u>0.037</u>	<u>0.036</u>
	DiffTraj	0.052	0.035	0.037
	TDDM	0.039	<u>0.037</u>	0.018
	TimeGAN	0.212	0.021	0.058
	TimeVAE	0.047	0.074	<u>0.005</u>
Length (↓)	COSCI-GAN	0.774	0.707	0.885
	Diffusion-TS	0.002	0.005	0.003
	DiffTraj	0.043	<u>0.012</u>	0.139
	TDDM	<u>0.004</u>	0.005	0.003
	TimeGAN	0.490	0.730	0.810
Pattern (↑)	TimeVAE	0.780	0.890	0.850
	COSCI-GAN	0.650	0.830	0.830
	Diffusion-TS	0.920	<u>0.900</u>	0.900
	DiffTraj	0.890	0.880	<u>0.910</u>
	TDDM	<u>0.910</u>	0.920	0.920

2210

2211

2212

2213

2214

2215

2216

2217

2218

2219

2220

Table 8: Ablation study of the effect of region size and spatial prior.

Measure	Model	Geolife	Porto	Cabspotting
TSTR (\downarrow)	TDDM	0.010 ± 0.005	0.009 ± 0.004	0.015 ± 0.008
	1x1 km	0.030 ± 0.014	0.021 ± 0.011	0.022 ± 0.011
	w/o spatial prior	0.009 ± 0.006	0.011 ± 0.005	0.012 ± 0.007
	w/o spatial prior + rejection	0.021 ± 0.009	0.007 ± 0.004	0.015 ± 0.008
KL($S \parallel R$) (\downarrow)	TDDM	0.202	0.437	0.265
	1x1 km	0.279	0.552	0.184
	w/o spatial prior	0.926	2.349	1.433
	w/o spatial prior + rejection	1.343	2.550	1.881
KL($R \parallel S$) (\downarrow)	TDDM	0.177	0.336	0.246
	1x1 km	0.206	0.579	0.168
	w/o spatial prior	0.858	1.314	1.123
	w/o spatial prior + rejection	0.890	1.478	1.388
KL _{sym} (\downarrow)	TDDM	0.190	0.386	0.255
	1x1 km	0.243	0.566	0.176
	w/o spatial prior	0.892	1.831	1.278
	w/o spatial prior + rejection	1.117	2.014	1.635
JS (\downarrow)	TDDM	0.040	0.079	0.056
	1x1 km	0.054	0.121	0.038
	w/o spatial prior	0.156	0.296	0.231
	w/o spatial prior + rejection	0.187	0.324	0.285
KL _{speed} (\downarrow)	TDDM	0.011	0.020	0.007
	1x1 km	1.097	0.528	0.123
	w/o spatial prior	0.478	0.124	0.368
	w/o spatial prior + rejection	0.406	0.215	0.644
Density (\downarrow)	TDDM	<u>0.023</u>	0.019	0.014
	1x1 km	0.021	<u>0.033</u>	0.011
	w/o spatial prior	0.099	0.043	0.060
	w/o spatial prior + rejection	0.086	0.046	0.055
Trip (\downarrow)	TDDM	0.039	0.037	0.018
	1x1 km	0.049	0.066	0.018
	w/o spatial prior	0.099	0.055	0.068
	w/o spatial prior + rejection	0.103	0.070	0.070
Length (\downarrow)	TDDM	0.004	0.005	0.003
	1x1 km	0.256	0.152	0.043
	w/o spatial prior	0.095	0.034	0.105
	w/o spatial prior + rejection	0.074	0.045	0.105
Pattern (\uparrow)	TDDM	0.910	0.920	0.920
	1x1 km	0.930	0.910	0.950
	w/o spatial prior	0.860	0.850	0.790
	w/o spatial prior + rejection	0.880	0.870	0.830

2263

2264

2265

2266

2267

2268
 2269
 2270
 2271
 2272
 2273
 2274
 2275
 2276
 2277

2278 Table 9: Map matching ablation study. Models trained without map matching during preprocessing.

2279	Measure	Model	Geolife	Porto	Cabspotting
2280	TSTR (\downarrow)	Diffusion-TS	0.015 ± 0.006	0.012 ± 0.005	0.020 ± 0.009
2281		DiffTraj	0.019 ± 0.005	0.007 ± 0.004	0.013 ± 0.007
2282		TDDM	0.013 ± 0.010	<u>0.010 ± 0.004</u>	0.011 ± 0.005
2283	$KL(S \parallel R) (\downarrow)$	Diffusion-TS	<u>1.196</u>	<u>2.401</u>	1.591
2284		DiffTraj	3.068	2.799	<u>1.623</u>
2285		TDDM	0.797	1.779	1.747
2286	$KL(R \parallel S) (\downarrow)$	Diffusion-TS	<u>1.038</u>	1.309	<u>1.263</u>
2287		DiffTraj	1.387	1.545	1.036
2288		TDDM	0.739	<u>1.451</u>	1.469
2289	$KL_{\text{sym}} (\downarrow)$	Diffusion-TS	<u>1.117</u>	<u>1.855</u>	<u>1.427</u>
2290		DiffTraj	2.227	2.172	1.329
2291		TDDM	0.768	1.615	1.608
2292	JS (\downarrow)	Diffusion-TS	<u>0.182</u>	<u>0.298</u>	<u>0.249</u>
2293		DiffTraj	0.319	0.341	0.238
2294		TDDM	0.138	0.259	0.264
2295	$KL_{\text{speed}} (\downarrow)$	Diffusion-TS	<u>0.058</u>	<u>0.018</u>	<u>0.042</u>
2296		DiffTraj	0.381	0.034	0.416
2297		TDDM	0.012	0.013	0.017
2298	Density (\downarrow)	Diffusion-TS	<u>0.051</u>	<u>0.025</u>	0.055
2299		DiffTraj	0.119	0.078	<u>0.043</u>
2300		TDDM	0.038	0.022	0.040
2301	Trip (\downarrow)	Diffusion-TS	<u>0.064</u>	<u>0.044</u>	0.065
2302		DiffTraj	0.127	0.087	0.051
2303		TDDM	0.052	0.037	<u>0.057</u>
2304	Length (\downarrow)	Diffusion-TS	<u>0.005</u>	<u>0.007</u>	0.005
2305		DiffTraj	0.139	0.062	0.286
2306		TDDM	0.003	0.004	<u>0.006</u>
2307	Pattern (\uparrow)	Diffusion-TS	<u>0.870</u>	0.890	0.790
2308		DiffTraj	0.800	0.810	<u>0.860</u>
2309		TDDM	0.910	<u>0.880</u>	0.880

2313
 2314
 2315
 2316
 2317
 2318
 2319
 2320
 2321

2322
 2323
 2324
 2325
 2326
 2327
 2328
 2329
 2330
 2331
 2332 Table 10: Robustness to map dropout. Models trained with random dropout of map information
 2333 during training.

Measure	Model	Geolife	Porto	Cabspotting
TSTR (↓)	TDDM	0.010 ± 0.005	0.009 ± 0.004	0.015 ± 0.008
	10% dropout	0.027 ± 0.014	<u>0.030 ± 0.016</u>	0.020 ± 0.011
	20% dropout	<u>0.025 ± 0.011</u>	0.033 ± 0.017	0.021 ± 0.010
KL($S \parallel R$) (↓)	TDDM	0.202	0.437	0.265
	10% dropout	0.246	<u>0.213</u>	0.214
	20% dropout	<u>0.243</u>	0.211	<u>0.215</u>
KL($R \parallel S$) (↓)	TDDM	0.177	0.336	0.246
	10% dropout	0.379	<u>0.200</u>	<u>0.253</u>
	20% dropout	<u>0.376</u>	0.197	0.256
KL _{sym} (↓)	TDDM	0.190	0.386	0.255
	10% dropout	0.313	<u>0.207</u>	0.234
	20% dropout	<u>0.310</u>	0.204	<u>0.235</u>
JS (↓)	TDDM	0.040	0.079	0.056
	10% dropout	<u>0.061</u>	0.045	0.049
	20% dropout	<u>0.061</u>	0.044	0.049
KL _{speed} (↓)	TDDM	0.011	0.020	0.007
	10% dropout	<u>4.442</u>	<u>2.422</u>	<u>0.837</u>
	20% dropout	4.523	3.096	1.279
Density (↓)	TDDM	<u>0.023</u>	<u>0.019</u>	<u>0.014</u>
	10% dropout	0.016	0.011	0.011
	20% dropout	0.016	0.011	0.011
Trip (↓)	TDDM	0.039	0.037	<u>0.018</u>
	10% dropout	0.032	0.034	0.017
	20% dropout	<u>0.033</u>	<u>0.035</u>	<u>0.018</u>
Length (↓)	TDDM	0.004	0.005	0.003
	10% dropout	<u>0.442</u>	<u>0.418</u>	<u>0.107</u>
	20% dropout	0.445	0.429	0.136
Pattern (↑)	TDDM	<u>0.910</u>	<u>0.920</u>	<u>0.920</u>
	10% dropout	0.920	0.930	0.950
	20% dropout	<u>0.910</u>	<u>0.920</u>	<u>0.950</u>

2376
 2377
 2378
 2379
 2380
 2381
 2382
 2383
 2384
 2385
 2386
 2387
 2388
 2389
 2390
 2391

Table 11: Generalization experiment.

Measure	Trained on	Geolife	Porto	Cabspotting
TSTR (↓)	100%	0.007 ± 0.004	0.006 ± 0.003	<u>0.016 ± 0.009</u>
	25%	<u>0.011 ± 0.006</u>	<u>0.010 ± 0.004</u>	0.014 ± 0.008
KL($S \parallel R$) (↓)	100%	0.199	0.448	0.266
	25%	<u>0.424</u>	<u>0.916</u>	<u>0.461</u>
KL($R \parallel S$) (↓)	100%	0.168	0.340	0.245
	25%	<u>0.373</u>	<u>0.629</u>	<u>0.403</u>
KL _{sym} (↓)	100%	0.184	0.394	0.256
	25%	<u>0.399</u>	<u>0.773</u>	<u>0.432</u>
JS (↓)	100%	0.039	0.080	0.056
	25%	<u>0.076</u>	<u>0.145</u>	<u>0.093</u>
KL _{speed} (↓)	100%	0.011	0.019	0.007
	25%	<u>0.046</u>	<u>0.027</u>	<u>0.230</u>
Density (↓)	100%	0.019	0.016	0.011
	25%	<u>0.025</u>	<u>0.020</u>	<u>0.014</u>
Trip (↓)	100%	0.032	0.037	0.013
	25%	<u>0.039</u>	<u>0.048</u>	<u>0.019</u>
Length (↓)	100%	0.003	0.005	0.002
	25%	<u>0.016</u>	<u>0.008</u>	<u>0.049</u>
Pattern (↑)	100%	0.930	0.930	0.960
	25%	<u>0.910</u>	<u>0.900</u>	<u>0.950</u>

2416
 2417
 2418
 2419
 2420
 2421
 2422
 2423
 2424
 2425
 2426
 2427
 2428
 2429

2430
 2431
 2432
 2433
 2434
 2435
 2436
 2437
 2438

2439 Table 12: City to city generalization experiment. The model is trained on one city and then is sampled
 2440 for a target city using the marginal distribution.

2441 2442 2443 Measure	2444 Trained on	2445 Sampled for		
		2446 Geolife	2447 Porto	2448 Cabspotting
2449 TSTR (\downarrow)	2450 Geolife	0.010 \pm 0.005	0.015 \pm 0.007	0.017 \pm 0.008
	2451 Porto	<u>0.011 \pm 0.006</u>	0.009 \pm 0.004	0.009 \pm 0.003
	2452 Cabspotting	0.012 \pm 0.006	<u>0.010 \pm 0.005</u>	0.015 \pm 0.008
2453 KL($S \parallel R$) (\downarrow)	2454 Geolife	0.202	1.168	0.637
	2455 Porto	<u>0.294</u>	0.437	<u>0.421</u>
	2456 Cabspotting	0.313	<u>0.907</u>	0.265
2457 KL($R \parallel S$) (\downarrow)	2458 Geolife	0.177	0.796	0.579
	2459 Porto	<u>0.232</u>	0.336	<u>0.394</u>
	2460 Cabspotting	0.258	<u>0.641</u>	0.246
2461 KL _{sym} (\downarrow)	2462 Geolife	0.190	0.982	0.608
	2463 Porto	<u>0.263</u>	0.386	<u>0.407</u>
	2464 Cabspotting	0.286	<u>0.774</u>	0.255
2465 JS (\downarrow)	2466 Geolife	0.040	0.175	0.123
	2467 Porto	<u>0.054</u>	0.079	<u>0.087</u>
	2468 Cabspotting	0.058	<u>0.145</u>	0.056
2469 KL _{speed} (\downarrow)	2470 Geolife	0.011	<u>0.505</u>	<u>0.139</u>
	2471 Porto	0.097	0.020	0.380
	2472 Cabspotting	<u>0.083</u>	0.703	0.007
2473 Density (\downarrow)	2474 Geolife	<u>0.023</u>	<u>0.021</u>	0.016
	2475 Porto	0.021	<u>0.019</u>	<u>0.015</u>
	2476 Cabspotting	0.026	0.019	0.014
2477 Trip (\downarrow)	2478 Geolife	0.039	<u>0.040</u>	<u>0.022</u>
	2479 Porto	0.047	0.037	0.026
	2480 Cabspotting	<u>0.042</u>	0.042	0.018
2481 Length (\downarrow)	2482 Geolife	0.004	<u>0.148</u>	<u>0.016</u>
	2483 Porto	0.028	0.005	0.092
	2484 Cabspotting	<u>0.024</u>	0.193	0.003
2485 Pattern (\uparrow)	2486 Geolife	0.910	0.920	0.930
	2487 Porto	0.930	0.920	0.930
	2488 Cabspotting	<u>0.920</u>	<u>0.910</u>	<u>0.920</u>

2489
 2490
 2491
 2492
 2493
 2494
 2495
 2496
 2497
 2498
 2499
 2500
 2501
 2502
 2503
 2504
 2505
 2506
 2507
 2508
 2509
 2510
 2511
 2512
 2513
 2514
 2515
 2516
 2517
 2518
 2519
 2520
 2521
 2522
 2523
 2524
 2525
 2526
 2527
 2528
 2529
 2530
 2531
 2532
 2533
 2534
 2535
 2536
 2537
 2538
 2539
 2540
 2541
 2542
 2543
 2544
 2545
 2546
 2547
 2548
 2549
 2550
 2551
 2552
 2553
 2554
 2555
 2556
 2557
 2558
 2559
 2560
 2561
 2562
 2563
 2564
 2565
 2566
 2567
 2568
 2569
 2570
 2571
 2572
 2573
 2574
 2575
 2576
 2577
 2578
 2579
 2580
 2581
 2582
 2583
 2584
 2585
 2586
 2587
 2588
 2589
 2590
 2591
 2592
 2593
 2594
 2595
 2596
 2597
 2598
 2599
 2600
 2601
 2602
 2603
 2604
 2605
 2606
 2607
 2608
 2609
 2610
 2611
 2612
 2613
 2614
 2615
 2616
 2617
 2618
 2619
 2620
 2621
 2622
 2623
 2624
 2625
 2626
 2627
 2628
 2629
 2630
 2631
 2632
 2633
 2634
 2635
 2636
 2637
 2638
 2639
 2640
 2641
 2642
 2643
 2644
 2645
 2646
 2647
 2648
 2649
 2650
 2651
 2652
 2653
 2654
 2655
 2656
 2657
 2658
 2659
 2660
 2661
 2662
 2663
 2664
 2665
 2666
 2667
 2668
 2669
 2670
 2671
 2672
 2673
 2674
 2675
 2676
 2677
 2678
 2679
 2680
 2681
 2682
 2683
 2684
 2685
 2686
 2687
 2688
 2689
 2690
 2691
 2692
 2693
 2694
 2695
 2696
 2697
 2698
 2699
 2700
 2701
 2702
 2703
 2704
 2705
 2706
 2707
 2708
 2709
 2710
 2711
 2712
 2713
 2714
 2715
 2716
 2717
 2718
 2719
 2720
 2721
 2722
 2723
 2724
 2725
 2726
 2727
 2728
 2729
 2730
 2731
 2732
 2733
 2734
 2735
 2736
 2737
 2738
 2739
 2740
 2741
 2742
 2743
 2744
 2745
 2746
 2747
 2748
 2749
 2750
 2751
 2752
 2753
 2754
 2755
 2756
 2757
 2758
 2759
 2760
 2761
 2762
 2763
 2764
 2765
 2766
 2767
 2768
 2769
 2770
 2771
 2772
 2773
 2774
 2775
 2776
 2777
 2778
 2779
 2780
 2781
 2782
 2783
 2784
 2785
 2786
 2787
 2788
 2789
 2790
 2791
 2792
 2793
 2794
 2795
 2796
 2797
 2798
 2799
 2800
 2801
 2802
 2803
 2804
 2805
 2806
 2807
 2808
 2809
 2810
 2811
 2812
 2813
 2814
 2815
 2816
 2817
 2818
 2819
 2820
 2821
 2822
 2823
 2824
 2825
 2826
 2827
 2828
 2829
 2830
 2831
 2832
 2833
 2834
 2835
 2836
 2837
 2838
 2839
 2840
 2841
 2842
 2843
 2844
 2845
 2846
 2847
 2848
 2849
 2850
 2851
 2852
 2853
 2854
 2855
 2856
 2857
 2858
 2859
 2860
 2861
 2862
 2863
 2864
 2865
 2866
 2867
 2868
 2869
 2870
 2871
 2872
 2873
 2874
 2875
 2876
 2877
 2878
 2879
 2880
 2881
 2882
 2883
 2884
 2885
 2886
 2887
 2888
 2889
 2890
 2891
 2892
 2893
 2894
 2895
 2896
 2897
 2898
 2899
 2900
 2901
 2902
 2903
 2904
 2905
 2906
 2907
 2908
 2909
 2910
 2911
 2912
 2913
 2914
 2915
 2916
 2917
 2918
 2919
 2920
 2921
 2922
 2923
 2924
 2925
 2926
 2927
 2928
 2929
 2930
 2931
 2932
 2933
 2934
 2935
 2936
 2937
 2938
 2939
 2940
 2941
 2942
 2943
 2944
 2945
 2946
 2947
 2948
 2949
 2950
 2951
 2952
 2953
 2954
 2955
 2956
 2957
 2958
 2959
 2960
 2961
 2962
 2963
 2964
 2965
 2966
 2967
 2968
 2969
 2970
 2971
 2972
 2973
 2974
 2975
 2976
 2977
 2978
 2979
 2980
 2981
 2982
 2983
 2984
 2985
 2986
 2987
 2988
 2989
 2990
 2991
 2992
 2993
 2994
 2995
 2996
 2997
 2998
 2999
 3000
 3001
 3002
 3003
 3004
 3005
 3006
 3007
 3008
 3009
 3010
 3011
 3012
 3013
 3014
 3015
 3016
 3017
 3018
 3019
 3020
 3021
 3022
 3023
 3024
 3025
 3026
 3027
 3028
 3029
 3030
 3031
 3032
 3033
 3034
 3035
 3036
 3037
 3038
 3039
 3040
 3041
 3042
 3043
 3044
 3045
 3046
 3047
 3048
 3049
 3050
 3051
 3052
 3053
 3054
 3055
 3056
 3057
 3058
 3059
 3060
 3061
 3062
 3063
 3064
 3065
 3066
 3067
 3068
 3069
 3070
 3071
 3072
 3073
 3074
 3075
 3076
 3077
 3078
 3079
 3080
 3081
 3082
 3083
 3084
 3085
 3086
 3087
 3088
 3089
 3090
 3091
 3092
 3093
 3094
 3095
 3096
 3097
 3098
 3099
 3100
 3101
 3102
 3103
 3104
 3105
 3106
 3107
 3108
 3109
 3110
 3111
 3112
 3113
 3114
 3115
 3116
 3117
 3118
 3119
 3120
 3121
 3122
 3123
 3124
 3125
 3126
 3127
 3128
 3129
 3130
 3131
 3132
 3133
 3134
 3135
 3136
 3137
 3138
 3139
 3140
 3141
 3142
 3143
 3144
 3145
 3146
 3147
 3148
 3149
 3150
 3151
 3152
 3153
 3154
 3155
 3156
 3157
 3158
 3159
 3160
 3161
 3162
 3163
 3164
 3165
 3166
 3167
 3168
 3169
 3170
 3171
 3172
 3173
 3174
 3175
 3176
 3177
 3178
 3179
 3180
 3181
 3182
 3183
 3184
 3185
 3186
 3187
 3188
 3189
 3190
 3191
 3192
 3193
 3194
 3195
 3196
 3197
 3198
 3199
 3200
 3201
 3202
 3203
 3204
 3205
 3206
 3207
 3208
 3209
 3210
 3211
 3212
 3213
 3214
 3215
 3216
 3217
 3218
 3219
 3220
 3221
 3222
 3223
 3224
 3225
 3226
 3227
 3228
 3229
 3230
 3231
 3232
 3233
 3234
 3235
 3236
 3237
 3238
 3239
 3240
 3241
 3242
 3243
 3244
 3245
 3246
 3247
 3248
 3249
 3250
 3251
 3252
 3253
 3254
 3255
 3256
 3257
 3258
 3259
 3260
 3261
 3262
 3263
 3264
 3265
 3266
 3267
 3268
 3269
 3270
 3271
 3272
 3273
 3274
 3275
 3276
 3277
 3278
 3279
 3280
 3281
 3282
 3283
 3284
 3285
 3286
 3287
 3288
 3289
 3290
 3291
 3292
 3293
 3294
 3295
 3296
 3297
 3298
 3299
 3300
 3301
 3302
 3303
 3304
 3305
 3306
 3307
 3308
 3309
 3310
 3311
 3312
 3313
 3314
 3315
 3316
 3317
 3318
 3319
 3320
 3321
 3322
 3323
 3324
 3325
 3326
 3327
 3328
 3329
 3330
 3331
 3332
 3333
 3334
 3335
 3336
 3337
 3338
 3339
 3340
 3341
 3342
 3343
 3344
 3345
 3346
 3347
 3348
 3349
 3350
 3351
 3352
 3353
 3354
 3355
 3356
 3357
 3358
 3359
 3360
 3361
 3362
 3363
 3364
 3365
 3366
 3367
 3368
 3369
 3370
 3371
 3372
 3373
 3374
 3375
 3376
 3377
 3378
 3379
 3380
 3381
 3382
 3383
 3384
 3385
 3386
 3387
 3388
 3389
 3390
 3391
 3392
 3393
 3394
 3395
 3396
 3397
 3398
 3399
 3400
 3401
 3402
 3403
 3404
 3405
 3406
 3407
 3408
 3409
 3410
 3411
 3412
 3413
 3414
 3415
 3416
 3417
 3418
 3419
 3420
 3421
 3422
 3423
 3424
 3425
 3426
 3427
 3428
 3429
 3430
 3431
 3432
 3433
 3434
 3435
 3436
 3437
 3438
 3439
 3440
 3441
 3442
 3443
 3444
 3445
 3446
 3447
 3448
 3449
 3450
 3451
 3452
 3453
 3454
 3455
 3456
 3457
 3458
 3459
 3460
 3461
 3462
 3463
 3464
 3465
 3466
 3467
 3468
 3469
 3470
 3471
 3472
 3473
 3474
 3475
 3476
 3477
 3478
 3479
 3480
 3481
 3482
 3483
 3484
 3485
 3486
 3487
 3488
 3489
 3490
 3491
 3492
 3493
 3494
 3495
 3496
 3497
 3498
 3499
 3500
 3501
 3502
 3503
 3504
 3505
 3506
 3507
 3508
 3509
 3510
 3511
 3512
 3513
 3514
 3515
 3516
 3517
 3518
 3519
 3520
 3521
 3522
 3523
 3524
 3525
 3526
 3527
 3528
 3529
 3530
 3531
 3532