

DutyTTE: Deciphering Uncertainty in Origin-Destination Travel Time Estimation

Xiaowei Mao^{1,2†}, Yan Lin^{3†}, Shengnan Guo^{1,2}, Yubin Chen^{1,2}, Xingyu Xian^{1,2},
Haomin Wen⁴, Qisen Xu^{1,2}, Youfang Lin^{1,2}, Huaiyu Wan^{1,2*}

¹School of Computer and Information Technology, Beijing Jiaotong University, Beijing, China

²Beijing Key Laboratory of Traffic Data Mining and Embodied Intelligence, Beijing, China

³Aalborg University, Aalborg, Denmark

⁴Carnegie Mellon University, Pittsburgh, USA

{maoxiaowei, guoshn, 20231244, xingyuxian, 23218133, yflin, hywan}@bjtu.edu.cn;
lyan@cs.aau.dk;

wenhaomin.whm@gmail.com

Abstract

Uncertainty quantification in travel time estimation (TTE) aims to estimate the confidence interval for travel time, given the origin (O), destination (D), and departure time (T). Accurately quantifying this uncertainty requires generating the most likely path and assessing travel time uncertainty along the path. This involves two main challenges: 1) Predicting a path that aligns with the ground truth, and 2) modeling the impact of travel time in each segment on overall uncertainty under varying conditions. We propose DutyTTE to address these challenges. For the first challenge, we introduce a deep reinforcement learning method to improve alignment between the predicted path and the ground truth, providing more accurate travel time information from road segments to improve TTE. For the second challenge, we propose a mixture of experts guided uncertainty quantification mechanism to better capture travel time uncertainty for each segment under varying contexts. Extensive experiments on two real-world datasets demonstrate the superiority of our proposed method.

Code — <https://github.com/maoxiaowei97/DutyTTE>

Introduction

Origin-destination (OD) travel time estimation (TTE) refers to estimating the travel time Δt from the origin (O) to the destination (D) starting from the departure time (T). In many real-world applications, providing just an average estimate of travel time is inadequate. A more reliable and informative approach includes quantifying travel time uncertainty. Specifically, we aim to estimate the confidence interval for travel time with a specified confidence level to quantify uncertainty, which is beneficial in many scenarios. For example, ride-hailing services can benefit from providing customers with lower and upper confidence bounds of travel time, allowing them to better plan their schedules (Liu et al. 2023). Moreover, understanding travel time uncertainty can help ride-hailing and logistics platforms improve decision-making effectiveness, such as in order dispatching and vehi-

*Corresponding author. †These authors contributed equally.
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

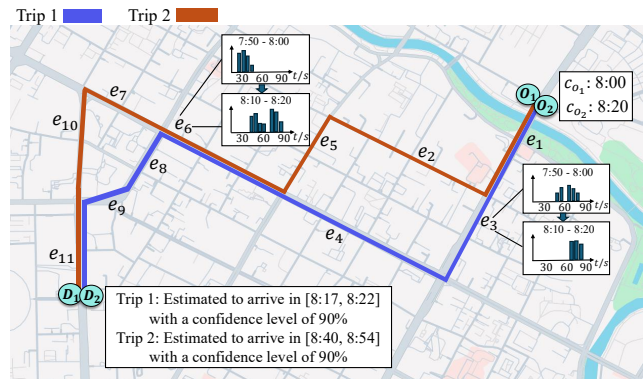


Figure 1: Motivation for DutyTTE: Each segment has varied travel time distributions, calculated using data from a period prior to the departure time.

cle routing (Xu et al. 2018; Fu et al. 2020). Existing studies for OD TTE (Yuan et al. 2020a; Lin et al. 2023; Wang et al. 2023) primarily focus on point estimates of travel time and have not adequately addressed uncertainty quantification.

Quantifying origin-destination travel time uncertainty involves two key tasks. First, a path between OD that matches the ground truth needs to be predicted, as the accuracy of the path is closely related to the accuracy of the quantified uncertainty in TTE. Second, the travel time uncertainty in each segment along the predicted path needs to be quantified so that the travel time for the whole trip can be aggregated (Lin et al. 2023; Wang et al. 2023). Effectively accomplishing these tasks faces several challenges.

First, predicting a path that aligns with the ground truth is challenging. Given an OD input, errors at multiple steps during path prediction can lead to significant divergence between the predicted and actual paths, especially for long paths with distant OD pairs. Such errors can accumulate, making the predicted path diverge from the actual path, which greatly hurts the accuracy of the quantified uncertainty in TTE. Traditional methods for path prediction typically involve modeling the weights of road segments, which are then integrated into search-based algorithms (Jain

et al. 2021; Tian et al. 2023) to generate a path step by step. More recent approaches (Shi et al. 2024; Mao et al. 2023b) use deep neural networks to learn transition patterns among roads and generate paths autoregressively. These methods face challenges in predicting paths that closely match the ground truth due to limited solutions for mitigating error accumulation during path prediction.

Second, effectively modeling the impact of travel time in each road segment on overall travel time uncertainty under varying conditions is challenging. Travel time on a segment can vary significantly between trips, contributing differently to the overall uncertainty. As illustrated in Figure 1, the travel time of e_6 in trip 1, which starts at 8:00, is relatively certain since the statistical times during 7:50–8:00 are concentrated around 30 seconds. However, during 8:10–8:20, travel times of e_6 become more varied, increasing the uncertainty for trip 2 which starts at 8:20. The varying segment-level travel time uncertainty, influenced by complex factors like traffic conditions and departure times, makes it difficult to accurately quantify overall path-level travel time uncertainty.

To address these challenges, we propose DutyTTE for Deciphering Uncertainty in origin-destination Travel Time Estimation. First, to reduce error accumulation in path prediction, we optimize path prediction by enhancing the overall alignment between the predicted path and the ground truth, rather than just maximizing the likelihood of each segment independently. We introduce a Deep Reinforcement Learning (DRL) approach to optimize objectives that measure the overall similarity between the predicted path and the ground truth. Second, we propose a Mixture of Experts guided Uncertainty Quantification mechanism (MoEUQ) to model the impact of travel time in each road segment on overall travel time uncertainty. This mechanism adaptively selects the most suitable experts to process each road segment, handling its travel time uncertainty and its impact in different contexts. This allows for better discrimination and capture of each road segment’s impact on overall travel time uncertainty.

Our contributions are summarized as follows:

- We propose an OD travel time uncertainty quantification method that accurately predicts paths and provides reliable confidence intervals.
- We introduce a DRL method to enhance the alignment between the predicted path and the ground truth, facilitating accurate travel time uncertainty quantification.
- We present a mixture of experts guided uncertainty quantification mechanism to model the impact of travel time in each road segment on overall travel time uncertainty.
- Extensive experiments on two real-world datasets show that our proposed method outperforms other solutions.

Related Work

Travel Time Estimation.

TTE solutions are classified into two types: for paths and for OD input.

Path-based TTE methods like Wide-Deep-Recurrent (WDR) (Wang, Fu, and Ye 2018a), DeepTTE (Wang et al.

2018), and ProbtTE (Liu et al. 2023) use deep learning to model spatio-temporal correlations in traveled segments. DeepGTT (Li et al. 2019) uses a variational model for TTE. GNNs and attention mechanism are used to capture spatial-temporal dependencies in road networks for TTE (Fang et al. 2020; Hong et al. 2020; Derrow-Pinion et al. 2021; Fu et al. 2020; Jin et al. 2022; Chen et al. 2022; Yuan, Li, and Bao 2022; James 2021; Ye et al. 2022; Liu et al. 2024; Mao et al. 2023a). ProbtTE (Xu, Wang, and Sun 2024) models link travel time with a Gaussian hierarchical model.

OD-based TTE methods, like TEMP (Wang et al. 2019), ST-NN (Jindal et al. 2017), MURAT (Li et al. 2018), DeepOD (Yuan et al. 2020b) model the correlations between ODT and travel times. DOT (Lin et al. 2023) develops a diffusion model for TTE using pixelated trajectories. MWSL-TTE (Wang et al. 2023) proposes to search a path first, and sums the estimated travel times of segments. However, these methods face limitations in modeling the complex correlations between OD pairs and travel time under varying traffic conditions, and they lack the capability to effectively model travel time uncertainty in road segments.

Uncertainty Quantification in Deep Learning. Methods for quantifying uncertainty in deep learning can be broadly classified into Bayesian and frequentist approaches (Wu et al. 2021). Bayesian methods (Izmailov et al. 2021; Kendall and Gal 2017) quantify uncertainty by assigning probabilistic distributions to model parameters, with variations in these parameters representing uncertainty. Non-Bayesian methods include ensemble-based approaches (Lakshminarayanan, Pritzel, and Blundell 2017; Liu et al. 2020), which model uncertainty by aggregating predictions from diverse models, and Brownian Motion-based techniques (Kong, Sun, and Zhang 2020), which simulate stochastic dynamics in the learning process to capture uncertainty. Techniques like MIS-Regression (Gneiting and Raftery 2007) and Quantile Regression (Gasthaus et al. 2019) estimate uncertainty intervals by modifying output layers without relying on likelihood functions. These uncertainty quantification methods can be integrated with TTE approaches to estimate travel time uncertainty (Zhu et al. 2022). However, such combinations fail to accurately predict paths and account for the varying uncertainty of travel times across segments under complex conditions.

Preliminaries

In this section, we provide the relevant definitions and formalize the problem of uncertainty quantification in origin-destination travel time estimation.

Definition 1. Road Network. A road network is defined as the aggregate of all road segments in a studied area or a city. It is modeled as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is a set of nodes v_i representing road intersections or segment ends, and \mathcal{E} is a set of edges e_i representing road segments. Each node in the network has a unique index.

Definition 2. Trips and Paths. A trajectory \mathcal{T} is a sequence of GPS points with timestamps: $\mathcal{T} = \langle (g_1, c_1), \dots, (g_{|\mathcal{T}|}, c_{|\mathcal{T}|}) \rangle$, where $g_i = (\text{lng}_i, \text{lat}_i)$, $i = 1, \dots, |\mathcal{T}|$ denotes i -th GPS point, and $|\mathcal{T}|$ denotes the total

number of GPS points in the trajectory. After map-matching, the trajectory of a trip is $\mathbf{x}_\tau = \{(v_1, c_1), \dots, (v_k, c_k)\}$, with the time index c monotonically increasing. The total travel time of the trip is $y = c_k - c_1$. Note that multiple GPS data points can be located on the same road segment. A path is defined as a sequence of nodes $\mathbf{x} = (v_1, \dots, v_{|\mathbf{x}|})$, where each pair of nodes is adjacent, i.e., $\forall i = 0, 1, \dots, |\mathbf{x}|, (v_i, v_{i+1}) \in \mathcal{E}$.

Problem Statement. We aim to learn a mapping function f_θ to estimate the upper and lower confidence bounds for travel time based on an ODT input, denoted as \hat{l} , and \hat{u} , which quantifies the uncertainty of the travel time. Such that $[\hat{l}, \hat{u}]$ is a confidence interval that covers the actual travel time τ with a confidence level of $1 - \rho$. We denote the ODT input as $\mathbf{q} = (g_o, g_d, c_o)$, where g_o, g_d , and c_o are the GPS points of the origin, destination, and departure time, respectively.

Methodology

Figure 2 illustrates the proposed method. This section discusses path prediction from a DRL perspective, policy-gradient optimization to align predicted paths with ground truth, and MoE-guided uncertainty quantification for travel time.

Path Prediction from the DRL Perspective

To mitigate error accumulation in path prediction, we aim to optimize overall alignment objectives between the predicted path $\hat{\mathbf{x}}$ and the ground truth \mathbf{x} . However, overall alignment measures like Longest Common Subsequence (LCS) (Bergroth, Hakonen, and Raita 2000) and Dynamic Time Warping (DTW) (Müller 2007) are non-differentiable when used as loss functions due to the arg max operation involved in obtaining the predicted path. To address this, we introduce a deep reinforcement learning (DRL) method to enhance path prediction by improving overall alignment.

Path prediction is modeled as a finite-horizon discounted Markov Decision Process (MDP), where an agent, over T discrete time steps, engages with the environment by making actions for path generation. An MDP is defined as $M = (S, A, P, R, s_0, \gamma)$, with S being the set of states and A being the set of possible actions. Next, $P : S \times A \times S \rightarrow \mathbb{R}_+$ is a transition probability function and $R : S \times A \rightarrow \mathbb{R}$ is a reward function. Finally, $s_0 : S \rightarrow \mathbb{R}_+$ is the initial state distribution and $\gamma \in [0, 1]$ is a discount factor.

Given a state s_t at time t , the agent uses the policy π_θ to make sequential decisions regarding the next node to generate. The agent then receives a reward r_t . Next, the reward is used to update the policy network. The objective is to learn the optimal parameters θ^* for the policy network to maximize the expected cumulative reward:

$\theta^* = \arg \max_{\theta} \mathbb{E}_{\pi_\theta} \left[\sum_{t=1}^T \gamma^t r_t \right]$, where γ controls the trade-offs between the importance of immediate and future rewards. Next, we describe agent, state, action, reward, and state transition in detail.

Agent. The agent takes actions according to the policy based on the states and receives rewards to learn an optimal policy that maximizes the cumulative reward. The actions

are determined by a policy network. The policy network is responsible for learning the mapping function f_θ from an ODT to a path. In our formulation, the policy network is mainly composed of a transformer (Vaswani et al. 2017) and a prediction head. The transformer generates the representation of the predicted path. Next, the prediction head takes the representation of the predicted path, traffic conditions, distance and direction to the destination as input, then it outputs the distribution of generating the next node.

State. The state at time step t , denoted as $s_t \in \mathcal{S}$, encodes context information necessary to derive a policy at each time step and is given by $s_t = (\mathbf{h}_t, \hat{\mathbf{x}}_{1:t-1}, w_t, dis_t, dir_t)$. Here, \mathbf{h}_t denotes the representation of the predicted path, $\hat{\mathbf{x}}_{1:t-1}$ is the previously predicted path, w_t denotes traffic conditions, dis_t is the distance from the current node to the destination, and dir_t is the direction from the current node to the destination at the t -th time step.

Action. An action $a_t \in \mathcal{A}_t$ represents choosing the next node to generate in the path. If there are T nodes in a path, the sequence $(a_1, a_2, \dots, a_T) \in \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_T$ represents a predicted path.

State Transition Probability. The transition probability function $P(s_{t+1}|s_t, a_t) : S \times A \times S \rightarrow \mathbb{R}_+$ quantifies the likelihood of transitioning from state s_t to state s_{t+1} given that action a_t is executed. In our problem, the environment is deterministic, meaning there is no uncertainty in the state transition, and s_{t+1} is uniquely determined.

Reward. The reward function, designed to enhance the alignment between the predicted path and the ground truth, is defined and described in detail in the following section.

Policy-Gradient Optimization for Path Prediction

Maximize the expected cumulative reward can be converted to a loss function defined as the negative expected reward of the predicted path: $\mathcal{L}_\theta = -\mathbb{E}_{\pi_\theta} [r(a_1, \dots, a_T)]$, where $\{a_1, \dots, a_T\} \sim \pi_\theta$ are sampled from the policy π_θ and $r(a_1, \dots, a_T)$ represents the reward. This reward is defined by overall alignment objectives, utilizing metrics such as DTW and LCS. DTW calculates the cost to align nodes of the predicted path with the ground truth in a way that minimizes the total distance. While LCS identifies the longest subsequence common to both the predicted path and the ground truth. Generating paths with lower DTW and higher LCS values indicates greater similarity to the ground truth. Formally, the reward function calculated using the predicted path $\hat{\mathbf{x}}$ and the ground truth \mathbf{x} is formulated as follows:

$$r(a_1, \dots, a_T) = (\omega * \text{LCS}(\hat{\mathbf{x}}, \mathbf{x}) - \beta * \text{DTW}(\hat{\mathbf{x}}, \mathbf{x})), \quad (1)$$

where ω , and β are hyper-parameters that control the scale of the reward. We can approximate the expected reward of the predicted path using N sampled paths from the policy π_{θ_a} . The derivative of loss function \mathcal{L}_θ is formulated as follows:

$$\nabla_{\pi_\theta} \mathcal{L}_{\pi_\theta} = -\frac{1}{N} \sum_{i=1}^N [\nabla_{\pi_\theta} \log \pi_\theta(a_{i,1}, \dots, a_{i,T}) \times r(a_{i,1}, \dots, a_{i,T})] \quad (2)$$

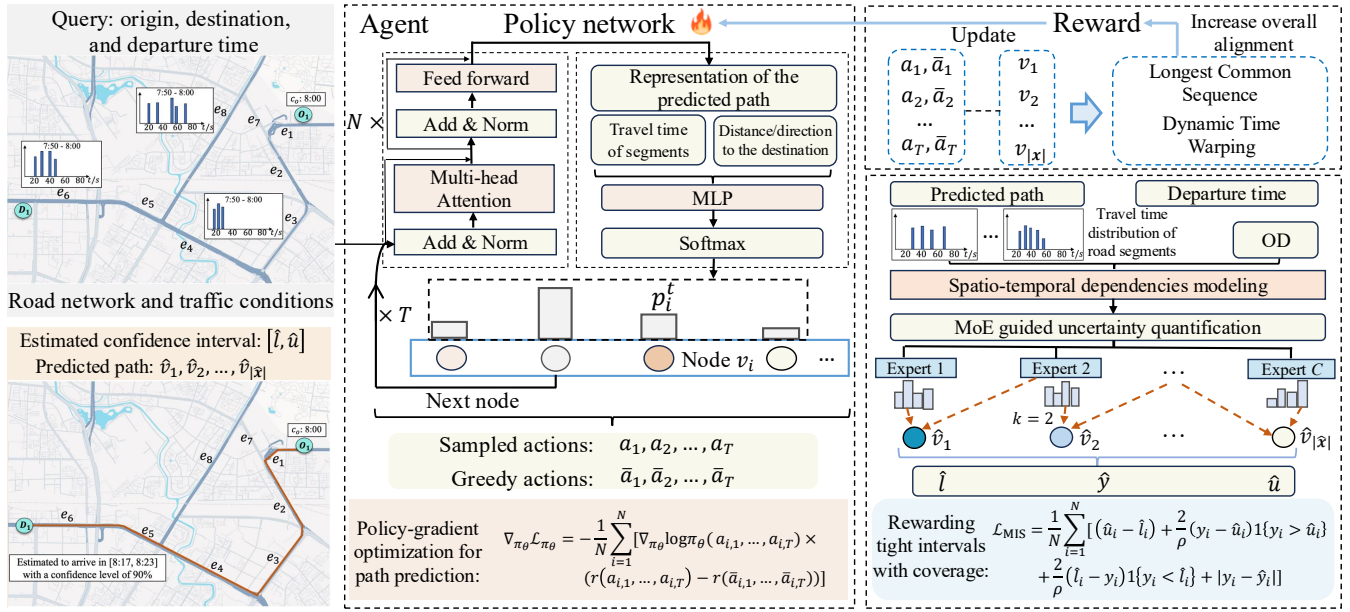


Figure 2: DutyTTE first learns to predict paths from an ODT input by optimizing objectives that measure the overall alignment. It then uses travel time information from segments in the predicted path to estimate travel time confidence intervals.

One limitation of this method is its high variance because it relies solely on samples when calculating the loss function during training. To address this, we use Self Critical Sequence Training (SCST) (Rennie et al. 2017). In SCST, the algorithm uses a baseline reward $r(\bar{a}_{i,1}, \dots, \bar{a}_{i,T})$ for the i -th sample, where $\bar{a}_{i,t}$ is the greedy selection of the model's prediction with frozen parameters. Thus, SCST compares samples from the model against its prediction with fixed parameters. This allows SCST to compare samples against predictions, reducing variance without changing the gradient's expectation (Rennie et al. 2017). We baseline the gradient using the reward from the arg max of the distribution when generating the next node at each step:

$$\nabla_{\pi_{\theta}} \mathcal{L}_{\pi_{\theta}} = -\frac{1}{N} \sum_{i=1}^N [\nabla_{\pi_{\theta}} \log \pi_{\theta}(a_{i,1}, \dots, a_{i,T}) \times (r(a_{i,1}, \dots, a_{i,T}) - r(\bar{a}_{i,1}, \dots, \bar{a}_{i,T}))] \quad (3)$$

where $\bar{a}_{i,t}$ is the arg max value of the output distribution for the next node from the t -th step of the i -th sample. If a sampled sequence is better than the model's greedy selection output, the probability of observing that sample is increased; if the sample is worse, the probability is decreased.

Inspired by (Hughes, Chang, and Zhang 2019; Paulus, Xiong, and Socher 2017), we also use Maximum-Likelihood Estimation (MLE) objective in path prediction to assist the policy learning algorithm to achieve better prediction. Formally, the joint loss function is formulated as follows:

$$\mathcal{L}_{\theta} = \gamma \mathcal{L}_{\pi_{\theta}} + \mathcal{L}_{\text{CE}}, \quad (4)$$

where γ is a hyper-parameter. \mathcal{L}_{CE} is defined by the MLE objective in path prediction (Shi et al. 2024).

MoE Guided Uncertainty Quantification

We propose a MoEUQ (Mixture of Experts for Uncertainty Quantification) model to capture the varying degrees of travel time uncertainty in each segment under different conditions. This segment-level uncertainty is then aggregated to estimate confidence intervals of overall travel time.

First, we take the discretized travel time distribution of segments in the predicted path as input, such as $\mathbf{u}_{e_j}^{c_o} = (d_1, f_1, \dots, d_m, f_m)$ for segment e_j . We select the top- m largest possible travel times in a period before departure time c_o . Here, d_m is the m -th largest travel time and f_m is the associated probability. We then concatenate $\mathbf{u}_{e_j}^{c_o}$ with learnable embeddings for the road segments and the time slice of the departure time, then encode them to obtain the embedding \mathbf{r}_j of the j -th segment. Next, we utilize a wide-deep-recurrent (Wang, Fu, and Ye 2018a; Mao et al. 2021) method to model the spatio-temporal dependencies among segments and the influence of other factors. The j -th encoded segment in the predicted path is denoted as \mathbf{r}'_j , which contains travel time information and aggregated spatio-temporal information of other segments in a path.

To better capture travel time uncertainty in each segment, we introduce the Mixture of Experts (MoE) approach (Shazeer et al. 2017). This method utilizes different combinations of experts to represent each segment based on its aggregated spatio-temporal information, effectively handling varying degrees of travel time uncertainty for the segment under different conditions. Specifically, given the j -th encoded segment \mathbf{r}'_j , the output of the MoE layer is:

$$\mathbf{r}''_j = \sum_{i=1}^C G(\mathbf{r}'_j) E_i(\mathbf{r}'_j), \quad (5)$$

where $E_i(\mathbf{r}'_j)$ denotes the output of the i -th expert network,

and $G(\mathbf{r}'_j)$ denotes the gating network given \mathbf{r}'_j . There are a total C expert networks with separate parameters, and the output of the gating network is a C -dimensional vector.

To prevent segments from being represented by the same set of experts and to explore better combinations of experts under varying conditions, we implement noisy top-k gating before applying the softmax function in the gating network. Each segment is routed to the most suitable experts, guided by the gating network, to distinguish travel time uncertainty in different conditions. The noisy top-k gating for the j -th segment in a predicted path is formulated as:

$$G(\mathbf{r}'_j) = \text{Softmax}(\text{TopK}(H(\mathbf{r}'_j), k)) \quad (6)$$

$$H(\mathbf{r}'_j)_i = (\mathbf{r}'_j \cdot \mathbf{W}_g)_i + \mathcal{N}(0, 1) \cdot \text{Softplus}((\mathbf{r}'_j \cdot \mathbf{W}_{\text{noise}})_i) \quad (7)$$

$$\text{TopK}(\mathbf{z}, k)_i = \begin{cases} z_i & \text{if } z_i \text{ is in the top } k \text{ elements of } \mathbf{z} \\ -\infty & \text{otherwise} \end{cases} \quad (8)$$

Next, to estimate the confidence intervals of travel time in the path, we sum along the sequence length dimension based on the output of the MoE layer. Then, we use separate prediction heads to estimate the lower and upper bounds of travel time, supervised by the Mean Interval Score (MIS) (Gneiting and Raftery 2007). For a confidence level of $1 - \rho$, the predicted upper and lower bounds of the travel time for the i -th sample given an ODT query \mathbf{q} are defined by $\hat{u}_i = \hat{y}_i + \hat{\sigma}_i^u$, and $\hat{l}_i = \hat{y}_i - \hat{\sigma}_i^l$, where \hat{u}_i and \hat{l}_i are the $(1 - \frac{\rho}{2})$ and $\frac{\rho}{2}$ quantiles for the $1 - \rho$ confidence interval and \hat{y}_i is the point estimation of travel time, respectively. The MIS loss is formulated as follows:

$$\mathcal{L}_{\text{MIS}} = \frac{1}{N} \sum_{i=1}^N \left[(\hat{u}_i - \hat{l}_i) + \frac{2}{\rho} (y_i - \hat{u}_i) \mathbf{1}\{y_i > \hat{u}_i\} + \frac{2}{\rho} (\hat{l}_i - y_i) \mathbf{1}\{y_i < \hat{l}_i\} + |y_i - \hat{y}_i| \right] \quad (9)$$

The MIS loss rewards narrower confidence intervals while encouraging them to include the actual arrival time.

Experiments

Comparison Methods

First, baselines for travel time uncertainty quantification include: **(1) T-WDR-MIS**: Using a transformer for path prediction, with WDR (Wang, Fu, and Ye 2018b) for TTE. It uses prediction heads to output confidence intervals, trained with MIS loss. **(2) MWSL-TTE** (Wang et al. 2023): Sums estimated travel time across all segments of the predicted path. Uncertainty is based on variance from a learned Gaussian distribution for each segment. **(3) DeepOD-MIS**: Uses DeepOD (Yuan et al. 2020b) with prediction heads for confidence intervals, trained with MIS loss. **(4) DOT-dropout**: Incorporates MC Dropout (Zhu and Laptev 2017) into DOT (Lin et al. 2023). **(5) DOT-MIS**: Equips DOT with prediction heads for confidence intervals, trained with MIS.

The desired confidence level of the confidence interval for travel time is set at 90% (Angelopoulos et al. 2022; Bates et al. 2021). To evaluate uncertainty quantification, we use Interval Width (IW) and Prediction Interval Coverage Probability (PICP) (Lawless and Fredette 2005; Zhou et al. 2021), where PICP is defined as the proportion of ground truth values covered by the estimated intervals.

Additional baselines for TTE: **(1) TEMP** (Wang et al. 2019): Estimate travel time using neighboring trips of a query. **(2) GBDT** (Chen and Guestrin 2016): A machine learning method. **(3) ST-NN**: (Jindal et al. 2017): Jointly predicts travel distance and time.

We also compare methods for path prediction: **(1) Dijkstra’s algorithm (DA)**: (Johnson 1973) Searches for the shortest path for given OD pairs. **(2) NMLR** (Jain et al. 2021): Learns segment weights and searches for paths with the largest weight at each step. **(3) Key Segment (KS)** (Tian et al. 2023): Detects a relay vertex for given OD pairs, predicting paths from origin to relay, then to the destination. **(4) Transformer**: Utilizes a transformer (Vaswani et al. 2017) for path prediction as implemented in (Shi et al. 2024). **(5) GDP** (Shi et al. 2024): uses a diffusion-based model to learn path distributions and predict paths via transformers.

Overall Comparison

Effectiveness of Travel Time Uncertainty Quantification T-WDR-MIS and MWSL-TTE outperform DeepOD-MIS in PICP by leveraging estimated travel time in the predicted path, highlighting the importance of predicting paths before quantifying uncertainty. T-WDR-MIS predicts paths autoregressively by maximizing the likelihood of each segment independently, while MWSL-TTE estimates transition weights between segments to predict paths. Both methods suffer from error accumulation in path prediction, limiting their ability to provide reliable path information for uncertainty quantification.

In contrast, DutyTTE enhances alignment with ground truth using DRL techniques to predict accurate paths while avoiding untraveled segment data. It employs MoEUQ to capture segment-level uncertainty in complex contexts by adaptively selecting experts to handle varying conditions. This enables DutyTTE to effectively learn the contribution of each segment’s travel time to overall uncertainty, achieving superior performance.

Effectiveness of Travel Time Estimation Table 3 presents the TTE results, where DutyTTE consistently outperforms other baselines across two datasets.

MWSL-TTE and T-WDR are more accurate than DOT and DeepOD by predicting potential paths and utilizing recent travel time information. T-WDR outperforms MWSL-TTE; while MWSL-TTE uses weak supervision for segment travel times, T-WDR leverages recent data in road segments and its wide-deep-recurrent architecture to capture correlations among influencing factors.

DutyTTE improves MAPE by 14.39% to 15.65% over the most competitive baseline. This is due to precise path prediction between OD pairs and a mixture of experts guided

Datasets	Chengdu					Xi'an				
	RMSE (↓)	MAE (s) (↓)	MAPE (%) (↓)	PICP (%) (↑)	IW (s)	RMSE (↓)	MAE (s) (↓)	MAPE (%) (↓)	PICP (%) (↑)	IW (s)
DOT-dropout	313.81	242.56	29.10	17.21	145.06	306.13	228.93	28.19	19.43	176.64
MWSL-TTE	307.18	239.62	29.86	73.82	782.07	304.89	227.76	27.58	75.03	802.74
DeepOD-MIS	320.51	253.73	31.87	70.49	837.68	319.74	251.06	30.62	68.26	856.26
DOT-MIS	315.42	242.79	30.60	71.84	806.15	307.15	233.26	28.70	73.25	819.51
T-WDR-MIS	306.89	221.52	27.92	84.57	829.39	299.13	215.61	26.81	84.93	840.28
DutyTTE (ours)	278.22	195.70	23.55	91.02	763.65	270.04	186.37	22.95	91.67	774.83
Improvement	9.34%	11.65%	15.65%	7.63%	-	9.72%	13.56%	14.39%	7.93%	-

Table 1: Performance of travel time uncertainty quantification.

Dataset	Chengdu	Xi'an
Time span	11.01–11.16, 2018	11.01–11.16, 2018
Sampling rate	3s	3s
Number of trajectories	1,613,355	1,951,585
average travel time	810.18s	770.32s
Area (width*height km ²)	14.38*9.98	10.14*8.87
Number of nodes	4599	5259
Number of edges	6737	7784

Table 2: Dataset statistics.

Datasets	Chengdu / Xi'an		
	RMSE (↓)	MAE (s) (↓)	MAPE (%) (↓)
TEMP	317.33/312.06	259.17/249.62	32.16/32.07
GBDT	315.91/318.10	247.03/245.83	31.03/30.12
STNN	324.62/336.56	268.67/264.95	33.28/32.93
DeepOD	320.51/319.74	253.73/251.06	31.87/30.62
MWSL-TTE	307.18/304.89	239.62/227.76	29.86/27.58
DOT	315.42/307.15	242.79/233.26	30.60/28.70
T-WDR	306.89/299.13	221.52/215.61	27.92/26.81
DutyTTE (ours)	278.22/270.04	195.70/186.37	23.55/22.95
Improvement	9.34%/9.72%	11.65%/13.56%	15.65%/14.39%

Table 3: Performance of travel time estimation.

travel time estimation method that effectively utilizes road segment travel time distributions under varying contexts.

Effectiveness of Path Prediction As shown in Table 4, KS surpasses DA and NMLR by identifying key segments and predicting paths based on constructed OD pairs with shorter distances. Both KS and NMLR rely on search-based methods that generate the next segment based solely on the learned weights among road segments. GDP models transition patterns in road networks using a diffusion model and uses a transformer to predict paths, yet still struggles to address error accumulation. DutyTTE addresses this by using a DRL method to enhance alignment between predicted and true paths, achieving superior performance. Figure 3 illustrates the reward curves for DutyTTE and a transformer implemented in GDP trained without policy loss (w/o policy). The higher rewards achieved by DutyTTE demonstrate enhanced path alignment resulting from DRL techniques.

Ablation Study We conduct an ablation study with the following DutyTTE variants: 1) w/o-M: removing the mixture of experts guided uncertainty quantification mechanism. 2) w/o-P: removing the policy loss and using only cross-entropy loss in path prediction.

Datasets	Chengdu		Xi'an	
	LCS (↑)	DTW (↓)	LCS (↑)	DTW (↓)
DA	3.74	0.103	4.62	0.093
NMLR	6.09	0.112	7.19	0.102
KS	12.03	0.039	12.82	0.036
Transformer	13.48	0.019	13.85	0.024
GDP	13.51	0.019	13.91	0.024
DutyTTE	14.52	0.017	14.64	0.021
Improvement	7.48%	10.53%	5.24%	12.50%

Table 4: Performance of path prediction.

Dataset	Chengdu				
	RMSE	MAE	MAPE	PICP	IW
w/o-M	282.41	198.47	24.32	88.63	781.52
w/o-P	302.93	219.32	27.28	87.15	836.09
DutyTTE	278.22	195.70	23.55	91.02	763.65

Table 5: Ablation study.




Figure 3: Reward curves during training.

We compare the performance of these variants with DutyTTE, and the experimental results are presented in Table 5. We observe the following: 1) The MoEUQ is crucial for accurate uncertainty quantification, as MoEUQ effectively learns the uncertainty of each segment under different contexts. 2) Integrating metrics that measure the overall similarity between predicted paths and the ground truth into the policy loss during training leads to improvements.

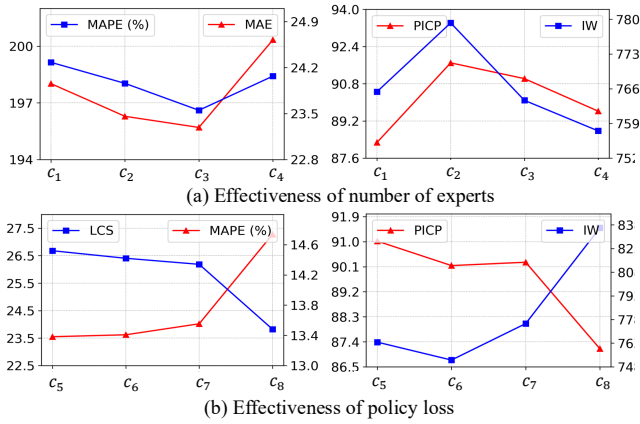


Figure 4: Effectiveness of hyper-parameters.

Hyper-Parameter Analysis We demonstrate the effectiveness of key hyper-parameters on testing sets. In Figure 4 (a), we set the value of k and the number of experts as follows: $c_1 : k = 1, C = 8$, $c_2 : k = 2, C = 8$, $c_3 : k = 4, C = 8$, $c_4 : k = 6, C = 8$. In Figure 4 (b), we configure the policy loss function with the following hyper-parameters: $c_5 : \omega = 1, \beta = 50, \gamma = 1$, $c_6 : \omega = 1, \beta = 100, \gamma = 1$, $c_7 : \omega = 1, \beta = 50, \gamma = 0.5$, $c_8 : \omega = 0, \beta = 0, \gamma = 0$. We observe the following results.

1) In Figure 4 (a), setting the value of k to 4 and the number of experts to 8 can lead to optimal performance. This can be attributed to the fact that having too few k and experts limits MoEUQ’s capability, as it can only handle a limited number of traffic contexts for evaluating travel time uncertainty in a segment. However, if the number of experts is too large, the model complexity also increases, complicating the training process and compromising performance.

2) In Figure 4 (b), applying the policy loss function with different parameter settings, such as c_5 , c_6 , and c_7 , results in improved performance compared to not using the policy loss function, as seen with c_8 . Among these configurations, c_5 stands out as a better option for more accurate path prediction compared to c_6 and c_7 . Achieving more accurate path predictions also leads to improvements in travel-time estimation and uncertainty quantification.

Analysis of the MoE Guided Uncertainty Quantification

The experimental results in Figures 5 and Figure 6 demonstrate the effectiveness of the MoE method in quantifying uncertainty for road segment travel times. In Figure 5, 2000 road segment samples under various traffic conditions were randomly selected, with travel time variance calculated from historical trajectories observed 10 minutes before trip departure. This variance reflects the uncertainty of travel time for each segment. The t-SNE visualization reduces high-dimensional embeddings of these segments into two dimensions for comparison. The results demonstrate that with MoE, clusters of road segments with low and high travel time variance (indicating low and high uncertainty) are more distinct compared to the scenario without MoE, where cluster boundaries are less defined. This underscores MoE’s abil-

Dataset	DA	DOT	T-WDR	MWSL	GDP	DutyTTE
Chengdu	0.0722	1.362	0.2675	0.1289	5.093	0.2706
Xi’an	0.0718	1.293	0.2649	0.1357	4.896	0.2682

Table 6: Execution time comparison (seconds per batch).




Figure 5: t-SNE Visualization for embeddings of road segments.




Figure 6: Expert activation distributions.

ity to better capture and differentiate travel time uncertainty under varying traffic conditions, improving overall uncertainty quantification. Additionally, Figure 6 shows the most frequently activated combination of experts for road segments with varying levels of uncertainty. The distribution of activated experts differs across uncertainty levels, indicating that the model dynamically employs diverse combinations of experts to effectively capture the uncertainty characteristics of road segments under different traffic conditions.

Efficiency Study We evaluate the average execution time per batch, consisting of 128 samples from Chengdu and Xi’an, for efficiency validation, as shown in Table 6. DutyTTE uses a DRL method for training; during prediction, the policy network allows for fast inferences. Given DutyTTE’s good performance, we see this as a promising trade-off between efficiency and effectiveness.

Conclusion

We introduce DutyTTE for quantifying uncertainty in OD travel times. Using a DRL approach, we optimize path prediction to align closely with the ground truth and enhance uncertainty quantification. To capture travel time uncertainty across road segments, we propose a mixture of experts mechanism that models the impact of segments in complex contexts. Experiments on two datasets validate DutyTTE’s effectiveness in path prediction, travel time uncertainty quantification.

Acknowledgments

This work was supported by the Fundamental Research Funds for the Central Universities (Grant No. 2023YJS030).

References

- Angelopoulos, A. N.; Kohli, A. P.; Bates, S.; Jordan, M.; Malik, J.; Alshaabi, T.; Upadhyayula, S.; and Romano, Y. 2022. Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In *International Conference on Machine Learning*, 717–730. PMLR.
- Bates, S.; Angelopoulos, A.; Lei, L.; Malik, J.; and Jordan, M. 2021. Distribution-free, risk-controlling prediction sets. *Journal of the ACM (JACM)*, 68(6): 1–34.
- Bergroth, L.; Hakonen, H.; and Raita, T. 2000. A survey of longest common subsequence algorithms. In *Proceedings Seventh International Symposium on String Processing and Information Retrieval. SPIRE 2000*, 39–48. IEEE.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *ACM SIGKDD*, 785–794.
- Chen, Z.; Xiao, X.; Gong, Y.-J.; Fang, J.; Ma, N.; Chai, H.; and Cao, Z. 2022. Interpreting trajectories from multiple views: A hierarchical self-attention network for estimating the time of arrival. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2771–2779.
- Derrow-Pinion, A.; She, J.; Wong, D.; Lange, O.; Hester, T.; Perez, L.; Nunkesser, M.; Lee, S.; Guo, X.; Wiltshire, B.; et al. 2021. Eta prediction with graph neural networks in google maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 3767–3776.
- Fang, X.; Huang, J.; Wang, F.; Zeng, L.; Liang, H.; and Wang, H. 2020. Constgat: Contextual spatial-temporal graph attention network for travel time estimation at baidu maps. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2697–2705.
- Fu, K.; Meng, F.; Ye, J.; and Wang, Z. 2020. CompactETA: A fast inference system for travel time prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3337–3345.
- Gasthaus, J.; Benidis, K.; Wang, Y.; Rangapuram, S. S.; Salinas, D.; Flunkert, V.; and Januschowski, T. 2019. Probabilistic forecasting with spline quantile function RNNs. In *The 22nd international conference on artificial intelligence and statistics*, 1901–1910. PMLR.
- Gneiting, T.; and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477): 359–378.
- Hong, H.; Lin, Y.; Yang, X.; Li, Z.; Fu, K.; Wang, Z.; Qie, X.; and Ye, J. 2020. HetETA: Heterogeneous information network embedding for estimating time of arrival. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2444–2454.
- Hughes, J. W.; Chang, K.-h.; and Zhang, R. 2019. Generating better search engine text advertisements with deep reinforcement learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2269–2277.
- Izmailov, P.; Vikram, S.; Hoffman, M. D.; and Wilson, A. G. G. 2021. What Are Bayesian Neural Network Posteriors Really Like? In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 4629–4640. PMLR.
- Jain, J.; Bagadia, V.; Manchanda, S.; and Ranu, S. 2021. Neuomlr: Robust & reliable route recommendation on road networks. *Advances in Neural Information Processing Systems*, 34: 22070–22082.
- James, J. 2021. Citywide estimation of travel time distributions with Bayesian deep graph learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(3): 2366–2378.
- Jin, G.; Wang, M.; Zhang, J.; Sha, H.; and Huang, J. 2022. STGNN-TTE: Travel time estimation via spatial-temporal graph neural network. *Future Generation Computer Systems*, 126: 70–81.
- Jindal, I.; Chen, X.; Nokleby, M.; Ye, J.; et al. 2017. A unified neural network approach for estimating travel time and distance for a taxi trip. *arXiv preprint arXiv:1710.04350*.
- Johnson, D. B. 1973. A note on Dijkstra’s shortest path algorithm. *Journal of the ACM*, 20(3): 385–388.
- Kendall, A.; and Gal, Y. 2017. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30.
- Kong, L.; Sun, J.; and Zhang, C. 2020. Sde-net: Equipping deep neural networks with uncertainty estimates. *arXiv preprint arXiv:2008.10546*.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- Lawless, J. F.; and Fredette, M. 2005. Frequentist prediction intervals and predictive distributions. *Biometrika*, 92(3): 529–542.
- Li, X.; Cong, G.; Sun, A.; and Cheng, Y. 2019. Learning travel time distributions with deep generative model. In *WWW*, 1017–1027.
- Li, Y.; Fu, K.; Wang, Z.; Shahabi, C.; Ye, J.; and Liu, Y. 2018. Multi-task representation learning for travel time estimation. In *ACM SIGKDD*, 1695–1704.
- Lin, Y.; Wan, H.; Hu, J.; Guo, S.; Yang, B.; Lin, Y.; and Jensen, C. S. 2023. Origin-Destination Travel Time Oracle for Map-Based Services. *Proc. ACM Manag. Data*, 1(3).
- Liu, H.; Jiang, W.; Liu, S.; and Chen, X. 2023. Uncertainty-Aware Probabilistic Travel Time Prediction for On-Demand Ride-Hailing at DiDi. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4516–4526.

- Liu, Y.; Qin, H.; Zhang, Z.; Pei, S.; Jiang, Z.; Feng, Z.; and Zhou, J. 2020. Probabilistic spatiotemporal wind speed forecasting based on a variational Bayesian deep learning model. *Applied Energy*, 260: 114259.
- Liu, Z.; Miao, H.; Zhao, Y.; Liu, C.; Zheng, K.; and Li, H. 2024. LightTR: A Lightweight Framework for Federated Trajectory Recovery. *arXiv preprint arXiv:2405.03409*.
- Mao, X.; Cai, T.; Peng, W.; and Wan, H. 2021. Estimated time of arrival prediction via modeling the spatial-temporal interactions between links and crosses. In *Proceedings of the 29th International Conference on Advances in Geographic Information Systems*, 658–661.
- Mao, X.; Wan, H.; Wen, H.; Wu, F.; Zheng, J.; Qiang, Y.; Guo, S.; Wu, L.; Hu, H.; and Lin, Y. 2023a. GMDNet: A graph-based mixture density network for estimating packages’ multimodal travel time distribution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4561–4568.
- Mao, X.; Wen, H.; Zhang, H.; Wan, H.; Wu, L.; Zheng, J.; Hu, H.; and Lin, Y. 2023b. Drl4route: A deep reinforcement learning framework for pick-up and delivery route prediction. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4628–4637.
- Müller, M. 2007. Dynamic time warping. *Information retrieval for music and motion*, 69–84.
- Paulus, R.; Xiong, C.; and Socher, R. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7008–7024.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Shi, D.; Tong, Y.; Zhou, Z.; Xu, K.; Wang, Z.; and Ye, J. 2024. Graph-constrained diffusion for end-to-end path planning. In *The Twelfth International Conference on Learning Representations*.
- Tian, W.; Shi, J.; Luo, S.; Li, H.; Xie, X.; and Zou, Y. 2023. Effective and efficient route planning using historical trajectories on road networks. *Proceedings of the VLDB Endowment*, 16(10): 2512–2524.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, D.; Zhang, J.; Cao, W.; Li, J.; and Zheng, Y. 2018. When will you arrive? Estimating travel time based on deep neural networks. In *AAAI*, volume 32, 2500–2507.
- Wang, H.; Tang, X.; Kuo, Y.-H.; Kifer, D.; and Li, Z. 2019. A simple baseline for travel time estimation using large-scale trip data. *ACM Trans. on Intelli. Sys. and Tech.*, 10(2): 1–22.
- Wang, H.; Zhang, Z.; Fan, Z.; Chen, J.; Zhang, L.; Shibasaki, R.; and Song, X. 2023. Multi-task weakly supervised learning for origin–destination travel time estimation. *IEEE Transactions on Knowledge and Data Engineering*, 35(11): 11628–11641.
- Wang, Z.; Fu, K.; and Ye, J. 2018a. Learning to estimate the travel time. In *ACM SIGKDD*, 858–866.
- Wang, Z.; Fu, K.; and Ye, J. 2018b. Learning to estimate the travel time. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 858–866.
- Wu, D.; Gao, L.; Chinazzi, M.; Xiong, X.; Vespignani, A.; Ma, Y.-A.; and Yu, R. 2021. Quantifying uncertainty in deep spatiotemporal forecasting. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 1841–1851.
- Xu, C.; Wang, Q.; and Sun, L. 2024. Link Representation Learning for Probabilistic Travel Time Estimation. *arXiv preprint arXiv:2407.05895*.
- Xu, Z.; Li, Z.; Guan, Q.; Zhang, D.; Li, Q.; Nan, J.; Liu, C.; Bian, W.; and Ye, J. 2018. Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 905–913.
- Ye, Y.; Zhu, Y.; Markos, C.; and James, J. 2022. CatETA: A categorical approximate approach for estimating time of arrival. *IEEE Transactions on Intelligent Transportation Systems*, 23(12): 24389–24400.
- Yuan, H.; Li, G.; and Bao, Z. 2022. Route travel time estimation on a road network revisited: Heterogeneity, proximity, periodicity and dynamicity. *Proceedings of the VLDB Endowment*, 16(3): 393–405.
- Yuan, H.; Li, G.; Bao, Z.; and Feng, L. 2020a. Effective Travel Time Estimation: When Historical Trajectories over Road Networks Matter. In *SIGMOD*, 2135–2149.
- Yuan, H.; Li, G.; Bao, Z.; and Feng, L. 2020b. Effective travel time estimation: When historical trajectories over road networks matter. In *Proceedings of the 2020 acm sigmod international conference on management of data*, 2135–2149.
- Zhou, Z.; Wang, Y.; Xie, X.; Qiao, L.; and Li, Y. 2021. Stuanet: Understanding uncertainty in spatiotemporal collective human mobility. In *Proceedings of the Web Conference 2021*, 1868–1879.
- Zhu, L.; and Laptev, N. 2017. Deep and confident prediction for time series at uber. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, 103–110. IEEE.
- Zhu, Y.; Ye, Y.; Liu, Y.; and James, J. 2022. Cross-area travel time uncertainty estimation from trajectory data: a federated learning approach. *IEEE Transactions on Intelligent Transportation Systems*, 23(12): 24966–24978.