# Model Agnostic Meta-Learning for traffic assignment on disrupted networks

Serio Agriesti*[1], Guido Cantelmo[1], and Francisco Pereira[1]

[1]Department of Technology, Management and Economics, Technical University of Denmark, Denmark

## Short summary

Accurate traffic flow estimation from origin-destination matrices is challenging when historical and future conditions diverge significantly — a common occurrence during network disruptions such as road closures, flooding, or public demonstrations. These are precisely the scenarios where reliable predictions matter most, yet standard machine-learning models struggle to quickly adapt to unseen graph structures with different demand patterns. This paper presents a meta-learning framework combined with a graph convolutional neural network (GNN) to address this limitation. Rather than requiring exhaustive training datasets covering all possible disruption scenarios, the architecture is trained to rapidly adapt to new network topologies and OD matrices with minimal data. Applied to simultaneous changes in road network structure and travel demand, the proposed model achieves an $R^2$ of approximately 0.85 on unseen network closures, significantly easing the burden on transport practitioners designing training datasets.
**Keywords**: Machine-Learning, Network disruption, Traffic Assignment.

## 1 Introduction

Graph Neural Networks have been widely adopted to surrogate traffic assignment models, as their structure strongly aligns with the nature of the transportation problem (T. Liu & Meidani, 2024, 2025; Hu & Xie, 2025; Lassen et al., 2025), a surrogate model being a machine-learning tool learning to map input-output relationships. In (T. Liu & Meidani, 2024) a graph neural network (GNN) is built, in which virtual links are added to improve message passing between key locations. In (T. Liu & Meidani, 2025) a GNN is built instead to perform the assignment of different classes of vehicles thanks to a dedicated attention mechanism. In (Hu & Xie, 2025), the performance of a GNN on different scenarios testing multiple values of demand and link capacities is assessed. The authors of Lassen et al. (2025) build a surrogate aimed at learning the input-output mapping of Stochastic User Equilibrium assignment instead. Still, all these works focus on building a surrogate model for a fixed network mirroring the existing infrastructure. Thus, the machine-learning models are extensively trained under the assumption that the graph will stay fixed or suffer limited changes (e.g. the closure of few links). The models will work for the conditions they have been trained for but their performance may degrade under network disruptions such as network maintenance (G. Liu et al., 2023), public demonstrations or disruptions such as flooding (Costa et al., 2024) and other extreme weather events. The combinatorial nature of the problem means that with the number of links, the number of possible network configurations increases by much higher orders of magnitude. Rather than training the surrogate over the full range of values in different scenarios, we propose a meta-learning architecture that results in fast adaptation of the neural network when changes in the graph arise, without a full re-training of the surrogate. Meta-learning is a methodological approach designed to initialize weights that generalize efficiently across structurally related tasks. Tasks may be regression, classification, or reinforcement learning problems (Finn et al., 2017). Meta-learning has already been successfully applied to transportation problems such as the macroscopic fundamental diagram estimation in data-scarce settings (Roark et al., 2025), trajectory (Wang et al., 2025) and demand prediction (Jiang et al., 2025) and travel time estimation (Fan et al., 2022). In this work, we embed a GNN model in a meta-learning framework (Finn et al., 2017) with the explicit aim of training it to achieve the best starting point for fast adaptation to different graph structures over the road network.

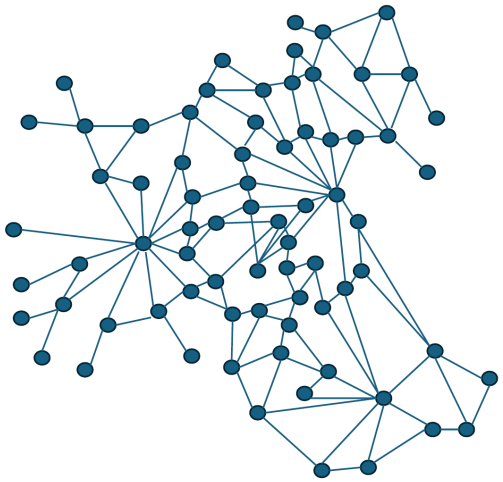


Figure 1: Road network as graph - 74 zones as nodes and 258 edges as links. Source: Stabler et al. (2026)

## 2 Methodology

In this section we describe the framework for model agnostic meta-learning (MAML) adapted from Finn et al. (2017) to fit the conditions of the problem at hand. The surrogate whose weights MAML computes is a gatedGCN, a type of GNN.

***The gatedGCN***

The surrogate is designed to predict flow volumes across a fixed road network, mapping the input origin-destination matrixes (ODs) to the flow on each link. Each node in the graph is a zone in the OD matrix, while the edges connecting them are considered links in the network (Fig. 1).
As node features we used: the number of edges converging, diverging from the node, the sum of the two and the flow generated towards all the other destinations. Capacity and a "present" flag were instead adopted as edge features. The "present" flag will allow the neural network to understand which edges are actually present in one specific configurations and which ones instead have been cut, to simulate disruptions once it is embedded in the MAML architecture. The gatedGCN architecture is composed by two linear layers (to lift node and edge features to latent space), 6 message passing layers representing the core of the surrogate's computations and 3 MLP layers to convert hidden representations into edge flows for each edge/link of the network. The message passing layers are mathematically defined by the following set of equations:

$$g = \sigma(W_e e_h + W_{dst} x_{dst} + W_{org} x_{org}) \tag{1}$$

$$m = (W_m x_{org}) \odot g \tag{2}$$

$$agg_{dst} = \frac{\sum m}{\sum g + \epsilon} \tag{3}$$

$$x' = W_{self} x_h + agg \tag{4}$$

where $W_e, W_{org}, W_{dst}, W_m, W_{self}$ are the weights the applied to the edge features $e_h$, to the node features $x_{org}, x_{dst}$, to the message $m$ that is passed between nodes and the hidden features $x_{org}$ of each node. $\epsilon$ is a small value to avoid dividing by zero and $x'$ is the updated representation of each node $x_h$ after each message passing layer. The gatedGCN implementation follow standard examples from literature, for a description of which the reader may refer to Bresson & Laurent (2018). A schematic representation is provided in Fig. 2.

***The MAML architecture***

The described gatedGCN architecture has been integrated into the MAML meta-learning framework. As noted in Section 1, MAML optimizes the initial weights of a surrogate model — the gatedGCN — to enable rapid adaptation to new experiments with limited data. To achieve this,

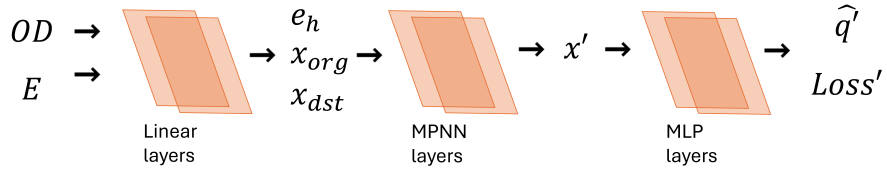



Figure 2: Schematic representation of one training epoch for the gatedGCN. $OD$ and $E$ are feature matrixes. $\hat{q}$ represents the estimated flows and $Loss' = SmoothL1Loss$ (SmoothL1Loss, 2026)

MAML simulates data scarcity during training and evaluates surrogate performance against the full dataset. In the inner-loop, the surrogate is exposed to a small data subset for a given task and optimizes its weights to minimize the corresponding loss. In the outer-loop, the surrogate is evaluated against the complete dataset. In our setting, each task is to produce accurate flow predictions on new network closures with scarce data (i.e., few OD matrix assignments). Accordingly, MAML trains the gatedGCN across multiple tasks (network closures), first presenting a limited number of OD matrix assignments (the support-set) and subsequently validating predictions on a larger set (the query-set). The full procedure is described in Alg. 1.

**Algorithm 1** Meta-training an adapted MAML for Few-Shot Supervised Learning

**Require:** $p(\mathcal{T})$: distribution over tasks
**Require:** $\alpha, \beta$: step size hyperparameters (learning rates)
  randomly initialize $\theta$
  **while** not done **do**
    Initialize outer loop
    Sample batch of tasks $\{\mathcal{T}_i\} \sim p(\mathcal{T})$
    **for all** $\mathcal{T}_i$ **do**
      Sample $K$ obs. $\mathcal{D}_i^{\text{support}} = \{x^{(j)}, y^{(j)}\}_{j=1}^{K}$ from $\mathcal{T}_i$
      Sample $M$ obs. $\mathcal{D}_i^{\text{query}} = \{x^{(j)}, y^{(j)}\}_{j=1}^{M}$ from $\mathcal{T}_i$
      Initialize task-specific parameters $\theta_i' \leftarrow \theta$
      **for** $N_{ite}$ **inner loop steps do**
        Compute K-specific loss:
        $\mathcal{L}_{\text{K}}(\mathcal{D}_i^{\text{support}}; \theta_i')$
        Average loss across $K$ samples:
        $\mathcal{L}_{\text{Task}} = \frac{1}{K} \sum \mathcal{L}_K$
        Update task-specific parameters:
        $\theta_i' \leftarrow \theta_i' - \alpha \nabla_{\theta_i'} \mathcal{L}_{\text{task}}(\mathcal{D}_i^{\text{support}}; \theta_i')$
      **end for**
      Compute query loss: $\mathcal{L}_{\text{task}}(\mathcal{D}_i^{\text{query}}; \theta_i')$
    **end for**
    Accumulate meta-gradient: $\nabla_\theta \sum_i \mathcal{L}_{\text{task}}(\mathcal{D}_i^{\text{query}}; \theta_i')$
    Update meta-parameters:
    $\theta \leftarrow \theta - \beta \nabla_\theta \sum_i \mathcal{L}_{\text{task}}(\mathcal{D}_i^{\text{query}}; \theta_i')$
  **end while**
  **return** $\theta$

In Alg.1 each OD assignment is one among $K$ examples while each set of closures corresponds to a task $\mathcal{T}$, for which $K$ OD assignments are carried out. MAML can thus train the meta-learned weights over different graphs (with the only constrain being that these share the same set of nodes).

## 3 The experimental setting

To generate the data needed for the experiments we used the traffic assignment tool Aequilibrae (Camargo, 2023). Through it, we carried out the Bi-conjugate Frank Wolfe assignment over the network of Eastern Massachussets (Camargo, 2023). The assignment has been run 24864 times to cover 336 different sets of closures and 74 different OD matrixes, the closures being characterized by random cuts ranging between 5% and 30% of the 258 edges. The 74 OD matrixes were instead generated as described in (Qurashi et al., 2021), to maintain spatial correlation while imposing a symmetric Gaussian distribution for which variations in trips range between 15% and 70%.

***Support- and query-set***

The generated data was structured to be used for the training of the gatedGCN through the MAML. This involved building two separated dataset: support and query. The support-set artificially imposes data scarcity, in our case simulating a situation in which few OD assignments are available for a specific network configuration. The query-set encompass a wider set of OD assignments (always excluding the ones already present in the support-set), so that the weights trained over support in the inner loop can be "validated" on a bigger range of configurations in the outer loop.

***Test set***

Once the MAML outer-loop meta-loss falls below an acceptable threshold, the resulting meta-weights are adopted by the gatedGCN and evaluated in their capacity to support rapid adaptation to unseen tasks (new closure configurations). To avoid overestimating meta-learning performance, these test tasks were kept entirely separated from meta-training (Alg. 1). To this end, 3 network closures and 25 OD matrices are held out from the original dataset prior to any training, yielding a test set of 3×25=75 assignments. Consistent with the few-shot adaptation objective (Finn et al., 2017), the test set is also partitioned into support and query subsets: the gatedGCN is permitted one round of inner-loop updates on the support data, simulating the limited OD observations available to the modeler upon encountering a new closure configuration. The resulting adapted weights are then assessed on their ability to generalize to the broader set of OD scenarios contained in the held-out query-set.

## 4 Results and discussion

In the experiments, the MAML framework and the GNN were trained to predict traffic flows in the test set: 3 unseen graphs (network closures 47, 248 and 293, randomly chosen) with 25 unseen ODs.

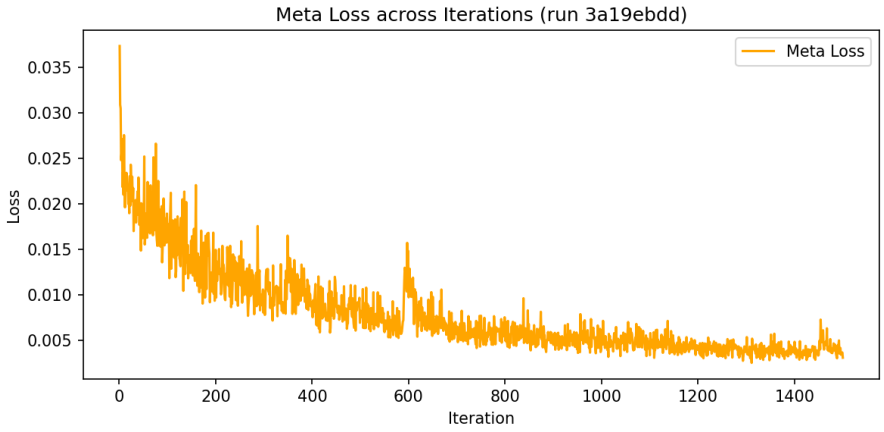



Figure 3: Meta-loss progression representing the improvements measured against the query-set in the outer loop, as the meta-training progresses

As it can be seen in Fig. 3, the meta-training results in a steady decrease of the loss function, as the MAML "shows" different tasks to the GNN, which then adapts its weights accordingly. The

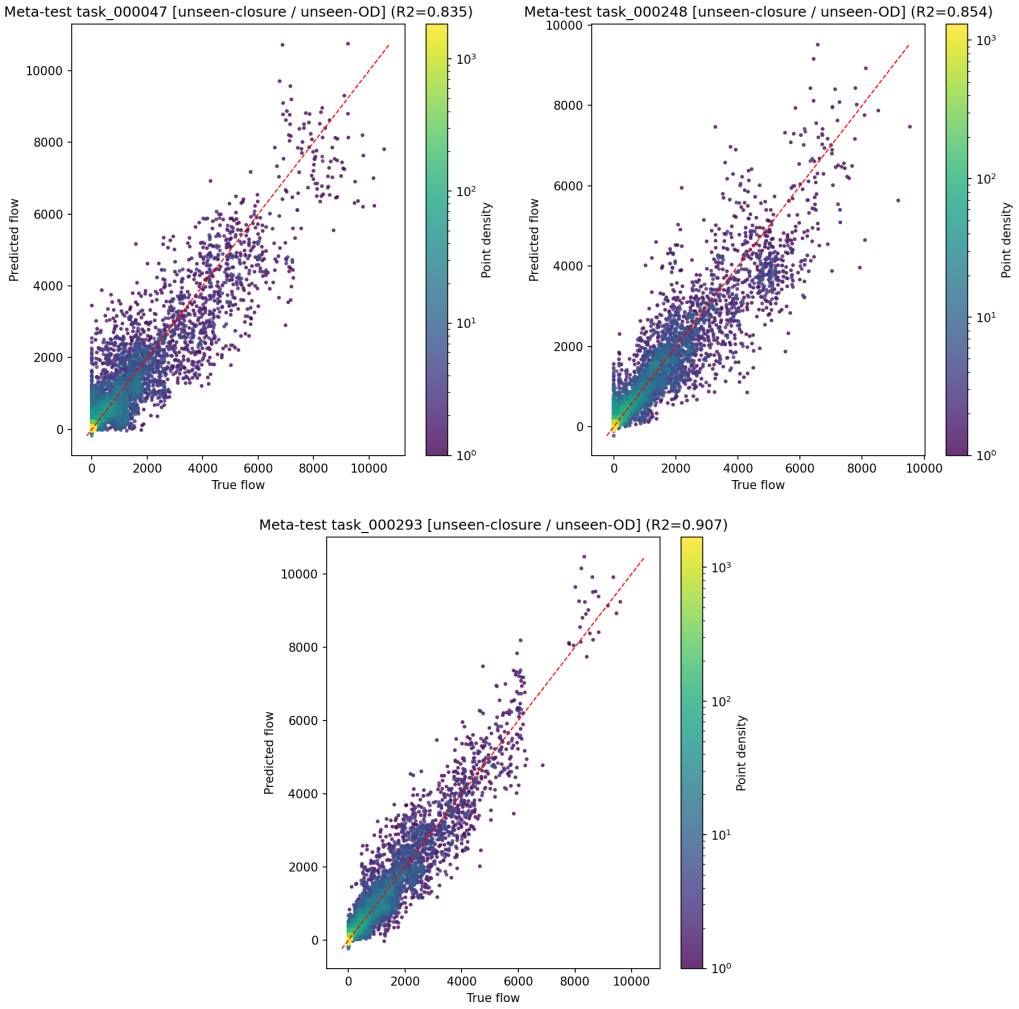



Figure 4: Scatter plots of predicted vs estimated flows

outer loss reaches a value of 0.003 on the normalized values of predicted and real flow for each link. From the trend it can be concluded that the MAML framework learns a good initializing starting point, as each meta-iteration exploits the same amount of inner loop iterations (Alg.1), the only difference being the starting weights themselves. Once the meta-loss is satisfactory, the weights are stored and applied to the test set. The gatedGCN applies the learned weights to predict flows arising from all the 25 unseen ODs on each set of network closures. The results are reported in Fig. 4

As the results show, the weights resulting from the MAML meta-training allow the gatedGCN to predict flows across new network configurations, proving that the meta-learning architecture learns how to perform the task of "predicting flows resulting from different ODs on new networks", rather than just mapping input to output. The experiment's settings are particularly demanding, as the gatedGCN must adapt rapidly not only to new closure configurations but also to OD demand matrixes entirely absent from training. This is most evident at $x = 0$: the MAML initialization enables the gatedGCN to correctly predict most of absent edges (reflected by the concentration of points at (0,0)). The observed $R^2$ range 0.835–0.907 suggests that closure topology may influence the model's few-shot adaptation capacity; nevertheless, predictions consistently align with the 45-degree line across all tasks, indicating robust prediction. To give a visual representation of the kind of network disruption the MAML is able to tackle, we report Fig. 5. As it can be noticed, the MAML+GNN architecture has been evaluated on tasks in which historical routing patterns are largely irrelevant for most OD pairs — scenarios where stakeholders are unlikely to possess sufficient demand data to retrain the gatedGCN, a relevant concern for medium-sized municipalities. By leveraging the proposed MAML framework, the gatedGCN maintains acceptable predictive performance without requiring retraining on new data or graph configurations. This translates directly into rapid response times, as reliable network flow estimates become available at the inference speeds characteristic of modern machine-learning solutions.

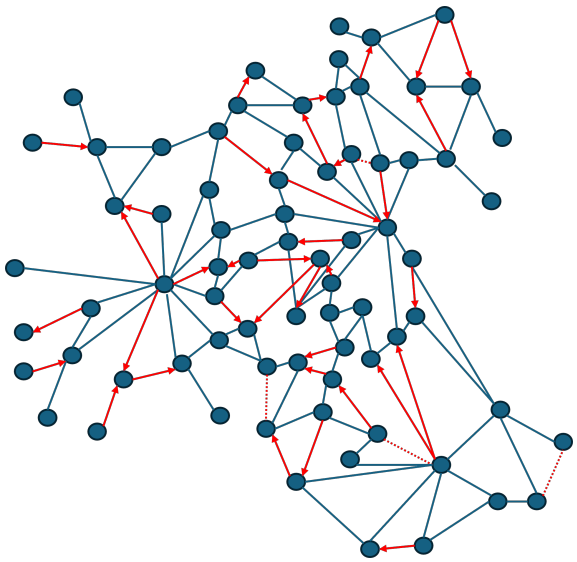


Figure 5: Network variation in task 47 - the red arrows signal the links that have become directed (one direction has been closed) while the dashed links represent closures in both directions.

## 5 Limitations and future research directions

We report a first application of the MAML framework on the problem of network disruptions, as the gatedGCN is meta-trained to learn the task of performing a traffic assignment on different sets of closures, with different available routes and travel times across ODs. One main limitation is that we did not allow for changes in the available nodes. This is a future research direction, so that data from other cities may be exploited in the meta-training. Varying both edges and nodes of the graph would indeed allow the modeler to gather data from multiple use cases suffering the same disruption and potentially improve the performance even further, while reducing the data requirements for the design of support and query-sets for the stakeholders from the considered city. Another limitation of the presented work is that the generation of OD matrixes is separated from the network closures with no direct relation between changes in demand and changes in the network structure. Adopting a more sophisticated simulation tool will improve the simulated datasets by directly accounting for this relation between demand and supply. We expect this improvement to actually ease the challenge for the MAML+GNN architecture, as it would allow the surrogate to learn an additional pattern that is currently absent in the meta-training dataset. Finally, the proposed framework should be tested on urban networks of varying scale and to alternative surrogate architectures, enabling systematic investigation of the effect of model design choices on meta-learning performance in traffic assignment settings.

## 6 Conclusions

This study presents a MAML-based framework for initializing a gatedGCN on traffic assignment tasks. Unlike prevailing approaches in the literature, the focus is not on forecasting but on learning how traffic flows distribute across varying graph topologies. Results demonstrate that embedding a standard gatedGCN within a MAML framework yields satisfactory performance despite the complexities introduced by structural graph changes. Although $R^2$ values range between 0.8 and 0.9, predictions consistently align with the 45-degree line across scatter plots, and the density of predicted versus true flow values indicates strong capacity to capture the effects of unseen closures on network flow patterns. The framework is well-suited for rapid adaptation during disruptive events that alter road network structure beyond what is captured by historical data, enabling deployment of a standard GNN without retraining.

## Acknowledgements

The work presented in this article is supported by Novo Nordisk Foundation grant NNF23OC0085356.

## References

Bresson, X., & Laurent, T. (2018). *Residual gated graph convnets.* (https://arxiv.org/abs/1711.07553)

Camargo, P. (2023). Aequilibrae and tradesman: current state of affairs of modelling with open-source software. In *In australasian transport research forum (atrf).*

Costa, M., Petersen, W., M., Vandervoort, A., Drews, M., Morrissey, K., & Pereira, C., F. (2024). *Climate adaptation with reinforcement learning: Experiments with flooding and transportation in copenhagen.* (https://arxiv.org/abs/2409.18574)

Fan, Y., Xu, J., Zhou, R., & Liu, C. (2022). Transportation-mode aware travel time estimation via meta-learning. In *In international conference on database systems for advanced applications* (p. 472-488).

Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *In international conference on machine learning* (p. 1126-1135).

Hu, X., & Xie, C. (2025). Use of graph attention networks for traffic assignment in a large number of network scenarios. *Transportation Research Part C: Emerging Technologies*, *171*, 104997.

Jiang, Z., Huang, A., Jiang, R., Chen, J., & Sekimoto, Y. (2025). Leveraging spatial–temporal heterogeneity and cross-mode interactions: A meta-learning approach for multimodal transportation demand prediction. *IEEE Transactions on Intelligent Transportation Systems*.

Lassen, O. B., Agriesti, S., Eldafrawi, M., Gammelli, D., Cantelmo, G., Gentile, G., & Pereira, F. C. (2025). Learning traffic flows: Graph neural networks for metamodelling traffic assignment. In *In 2025 9th international conference on models and technologies for intelligent transportation systems (mt-its)* (p. 1-8).

Liu, G., Zhang, X., Qian, Z., Chen, L., & Bi, Y. (2023). Life cycle assessment of road network infrastructure maintenance phase while considering traffic operation and environmental impact. *Journal of Cleaner Production*, *422*, 138607.

Liu, T., & Meidani, H. (2024). End-to-end heterogeneous graph neural networks for traffic assignment. *Transportation Research Part C: Emerging Technologies*, *165*, 104695.

Liu, T., & Meidani, H. (2025). Multi-class traffic assignment using multi-view heterogeneous graph attention networks. *Expert Systems with Applications*, *286*, 128072.

PyTorch. (2026). ([Online:] https://docs.pytorch.org/docs/stable/generated/torch.nn.SmoothL1Loss.html)

Qurashi, M., Lu, Q., Cantelmo, G., & Antoniou, C. (2021). Pc-spsa: Exploration and assessment of different historical data–set generation methods for enhanced dta model calibration. In *In 9th symposium of the european association for research in transportation (heart2020).*

Roark, A., Agriesti, S., Pereira, F. C., & Cantelmo, G. (2025). *Learning to learn the macroscopic fundamental diagram using physics-informed and meta machine learning techniques.* (arXiv:2508.14137)

Stabler, B., Bar-Gera, H., & Sall, E. (2026). *Transportation networks for research core team. transportation networks for research.* ([Online:] https://github.com/bstabler/TransportationNetworks)

Wang, C., Zhao, F., Luo, H., Sun, P. Z., & Fang, Y. (2025). Cross-transportation-mode knowledge transfer for trajectory recovery with meta learning. *IEEE Transactions on Intelligent Transportation Systems*.

***Appendix***

***Hyperparameters***

Here we report the hyperparameters we adopted to ensure reproducibility. These are listed in Table 1 for the gatedGCN and in Table 2 for the MAML architecture.

| | |
|---|---|
| Epochs | 100 |
| Batch size | 128 |
| Hidden dimensions | 192 |
| Message passing layers | 6 |
| Dropout | 0.01 |
| Learning rate | 0.02 |

Table 1: Hyperparameters adopted for the gatedGCN

| | |
|---|---|
| Meta-iterations | 1500 |
| K | 4 |
| M | 25 |
| Inner loop steps | 5 |
| $\mathcal{T}_i$ batch | 7 |
| Learning rate | 0.055 |

Table 2: Hyperparameters adopted for the meta-training of the MAML architecture with M being the number of ODs included in the query-set for each task $\mathcal{T}_i$