Estimating Accessibility of new Demand-Responsive Transport Deployments via Transfer Learning

Mohamed Ourahou[1,3,4*], Andrea Araldo[1], Mohamed Abouelela[2], Louis Zigrand[3], Aline Carneiro Viana[4], and Constantinos Antoniou[2]

[1]Institut Polytechnique de Paris, Palaiseau, France
[2]Chair of Transport Systems Engineering, Technical University of Munich, Munich, Germany
[3]Padam Mobility - Siemens Group, Paris, France
[4]Inria, Palaiseau, France
[*]mohamed.ourahou@telecom-sudparis.eu

ABSTRACT

Demand-Responsive Transport (DRT) is increasingly introduced to improve mobility in areas where conventional Public Transport (PT) offers limited frequency or coverage. Given the high operational costs of DRT, it is crucial for its further adoption, to rigorously assess its benefits, considering not only traditional operational metrics such as waiting and travel times, but also accessibility. Accessibility measures the ease of reaching surrounding opportunities such as jobs, healthcare, or education. However, data-driven accessibility computation for integrated DRT-PT services remains challenging in newly deployed territories, where historical operational data are unavailable or scarce.

This paper proposes a Domain-Adaptive Multi-Task Learning (DAMTL) framework to learn job accessibility patterns from operational data of different "source" territories, in which DRT is already deployed. Through adversarial domain adaptation, such a knowledge is transferred to a target territory with limited data.

We apply the framework to three French metropolitan areas: Lille and Strasbourg as source territories, and Orléans as the target. Results show that *with approximately 5 weeks of DRT operation in the target territory, the accessibility estimates issued by our method approximate well (70.1%) the estimates achievable over $\sim$ 6 months.*

These results show that our method can be successfully employed in the initial phases of its deployment in a new territory (e.g., in a pilot), to estimate, from the beginning, the benefits of DRT to the accessibility of a territory. Such early accessibility estimation can be used to confirm the pertinence of a certain DRT deployment or to diagnose some undesired accessibility distribution outcome and correct it via appropriate service planning intervention.

Keywords: Accessibility; demand responsive transport; transfer learning, mobility on-demand; public transport; travel time variability

1 Introduction and Related Work

Transport networks are fundamental to urban life, enabling individuals to reach employment, healthcare, education, and other essential opportunities. The potential to reach such places is formalised under the concept of location-based accessibility [Hansen, 1959, Miller, 2020]. In dense urban areas, Public Transport (e.g., fixed-route buses, trams, metros, and trains) generally provides effective accessibility [Pucher, 2004]. However, in low-density or peripheral areas, PT coverage is inherently limited by economic viability constraints, forcing operators to reduce service frequency and coverage [Calabrò et al., 2023, Currie and Fournier, 2020]. This encourages car dependence and restricting access for transport-disadvantaged populations [Saeidizand et al., 2022, Lucas, 2012].

Accessibility is commonly quantified using cumulative opportunity indicators, which count the number of opportunities reachable within a given travel time budget [Miller, 2020, Conway et al., 2018]. For PT, these indicators are typically computed using Time-Expanded Graphs (TEGs) derived from GTFS data, where shortest-path algorithms determine earliest arrival times [Farber et al., 2014, Li et al., 2020, Fransen et al., 2019]. However, such graph-based methods are not directly applicable to Demand-Responsive Transport (DRT), whose routes and schedules are dynamically constructed in response to passenger requests.

DRT has emerged as a promising PT complement in areas with sparse or heterogeneous demand [Viergutz and Schmidt, 2019, Alonso-González et al., 2018], bridging first- and last-mile gaps and connecting peripheral areas to major PT corridors. Despite its potential, DRT deployment remains costly [Calabrò et al., 2023, Currie and Fournier, 2020], making rigorous

accessibility assessment a prerequisite for justified investment decisions. DRT travel times are inherently variable: beyond traffic conditions, stochastic detours driven by concurrent passenger requests render travel times difficult to predict a priori [Fielbaum and Alonso-Mora, 2020], making a single deterministic estimate insufficient for reliable accessibility computation.

Estimating travel demand is a prerequisite for DRT accessibility computation, as it determines which origin-destination (OD) pairs are likely to generate trips. Data-driven approaches range from zone-level aggregated demand models [Hua et al., 2024, Ke et al., 2017] to probabilistic frameworks [Peng et al., 2025], but typically require large volumes of historical data from the target territory. The spatial transferability of demand models has received limited attention: Roy et al. [2025] show that transferred models can approximate local ride-hailing demand with reduced data requirements, but their work focuses on zone-level demand volume and does not extend to multimodal accessibility estimation.

Several studies have computed accessibility in flexible mobility systems. Diepolder et al. [2025] model DRT as a feeder service within a PT TEG, estimating travel times via kriging interpolation and assuming fixed-headway vehicle availability, without explicit demand modelling. Ourahou et al. [2025] extend this to many-to-many DRT configurations with probabilistic travel time estimation, restricting DRT connections to historically observed pick-up and drop-off locations as a demand proxy. Both methods rely on sufficient historical DRT data from the target territory, limiting applicability when the service is newly deployed or data are scarce — a limitation shared more broadly by data-driven accessibility approaches for PT [Bimpou and Ferguson, 2021, Braga et al., 2023] and flexible modes [Jiang et al., 2018].

To the best of our knowledge, no existing work has addressed accessibility estimation by DRT in territories where historical operational data are unavailable or scarce. This gap is practically significant: transport authorities increasingly need to anticipate DRT accessibility impacts prior to or in the early stages of deployment. This raises a practically and policy-relevant question: ***How much local operational data is needed to obtain a reliable estimate of accessibility by DRT in a new territory?*** This study addresses this gap by proposing a Domain-Adaptive Multi-Task Learning (DAMTL) framework that jointly estimates trip occurrence likelihood between OD pairs and associated travel time distributions, transfers the learned knowledge from source territories through adversarial domain adaptation, and integrates the resulting estimates into a multimodal TEG to compute cumulative job accessibility indicators. Moreover, the framework is evaluated under zero-shot and multiple few-shot transfer regimes.

2 Methodology and Data

This section presents a comprehensive methodology for estimating the accessibility provided by integrated DRT-PT services in new territories with limited or no historical data, using a transfer learning approach. It describes the methods, datasets, and case studies used to predict travel time distributions and derive accessibility indicators.

2.1 Accessibility computation

We measure accessibility to jobs using the cumulative accessibility indicator [Miller, 2020, Conway et al., 2018], which counts the total number of jobs reachable within a given travel time budget τ. This indicator is chosen for its simplicity, interpretability, and ease of communication.

The study area is divided into a regular hexagonal tessellation, with centroids denoted by the set $\mathcal{C} \subseteq \mathbb{R}^2$ as in Ourahou et al. [2025], Biazzo et al. [2019], Diepolder et al. [2023]. Each centroid serves as both an origin and a destination for travel time computation across the integrated DRT-PT network. The accessibility score for each hexagon is expressed as:

$$A_u(t) = \sum_{v \in \mathcal{C}} O_v \cdot f(T_{uv}(t)), \quad f(T_{uv}(t)) = \begin{cases} 1 & \text{if } T_{uv}(t) \leq \tau \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where $A_u(t)$ is the accessibility of hexagon u at departure time t, O_v is the number of jobs in hexagon v, $T_{uv}(t)$ is the shortest travel time from u to v departing at t using the integrated DRT-PT service, τ is the travel time threshold, and $f(\cdot)$ is the impedance function. While various impedance functions exist, we adopt this binary form as it enables intuitive interpretation of accessibility results and can be easily adapted or replaced as required.

2.2 Domain-Adaptive Multi-Task Learning Framework (DAMTL)

As established in Section 1, DRT's dynamic routing introduces time-dependent detours and substantial variability in travel times, making standard graph-based accessibility approaches inapplicable. Moreover, data-driven alternatives remain limited in data-scarce environments, particularly when transferring to new territories where no prior DRT data exist. To address these challenges, this paper proposes a DAMTL framework (Fig. 1) that jointly estimates (i) OD pair activation — used as a proxy for trip occurrence — and (ii) the travel time distribution of the corresponding trips. The framework leverages data from multiple source territories and employs adversarial domain adaptation to transfer learned representations to a target territory with limited

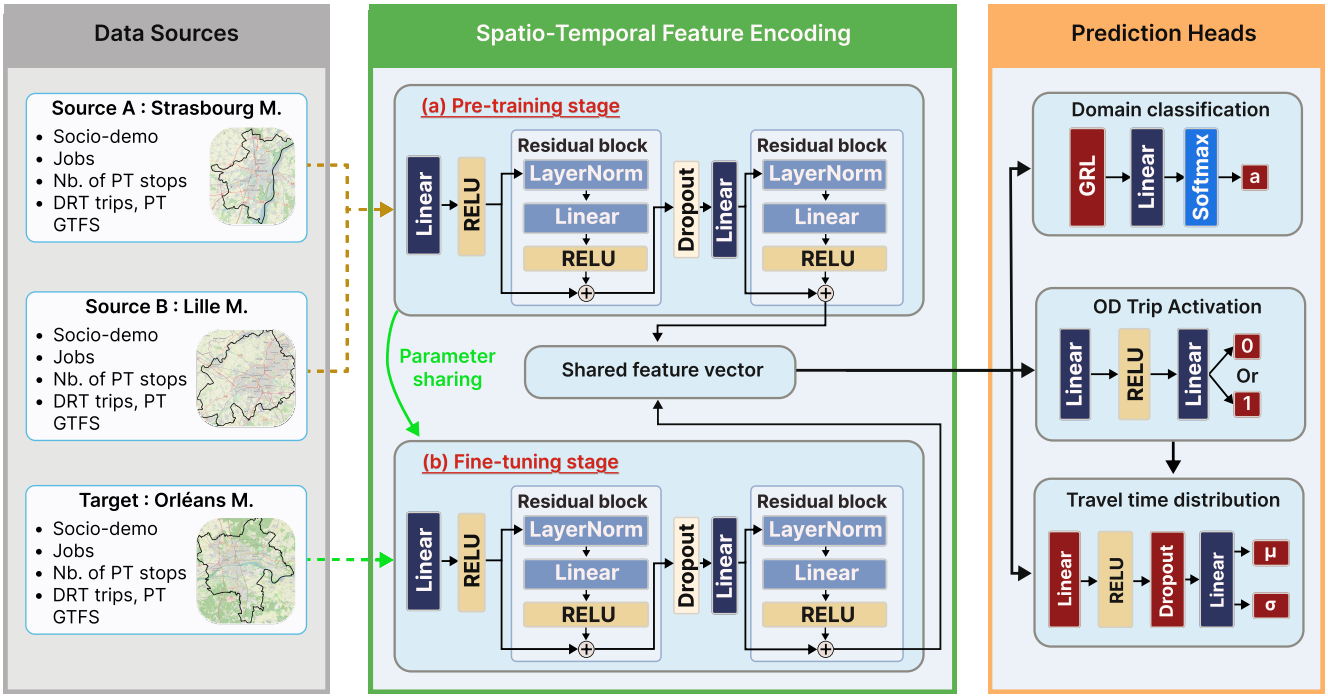


Figure 1. Domain-Adaptive Multi-Task Learning Framework (DAMTL) for joint OD activation and probabilistic travel time estimation across territories.

or no historical data. This enables robust estimation of DRT service characteristics, which are subsequently integrated into a multimodal DRT-PT accessibility analysis. The following subsections describe the components of the framework.

2.2.1 Data Sources

The DAMTL framework relies on four primary sources of data, retrieved for different territories.

- **Real-world Historical DRT trip data** are retrieved from [Padam Mobility, 2022]. Each trip is defined by departure and arrival geographic coordinates and times, providing the empirical foundation for demand characterisation and travel time estimation. For the OD activation head (Section 2.2.4), each recorded trip constitutes a positive sample ($y = 1$). For each positive sample, ten negative samples ($y = 0$) are generated by randomly drawing alternative OD hexagon pairs within the same territory, preserving departure time features. This 10:1 ratio reflects the real-world sparsity of active OD pairs while maintaining manageable class imbalance.
- **Sociodemographic data** for 2019 are obtained from INSEE [INSEE, 2025] at the IRIS level[1], including total population, share of residents aged over 65, median income per Consumption Unit, share of households below the poverty line, and commuting modes of active workers aged 15–64. All indicators are remapped onto a uniform hexagonal grid (2 km side length) using area-weighted redistribution: count-based variables are allocated proportionally to hexagon-IRIS overlap area, while rate-based indicators are first converted to absolute counts, redistributed, then recomputed as rates. Hexagons with fewer than 6 residents per km^2 are excluded.
- **Job counts** per IRIS zone are also obtained from INSEE [INSEE, 2025] and remapped to hexagons using the same area-weighting procedure. Job accessibility is prioritised as a target measure given its central role in daily mobility.
- **Road network data** are obtained from OpenStreetMap, and **PT schedules** from open GTFS feeds.

2.2.2 Spatio-Temporal Feature Encoding

The foundation of the DAMTL framework is a shared spatio-temporal encoder that maps heterogeneous, hexagon-level input features into a compact, domain-invariant representation. For each OD pair, the input feature vector is constructed by concatenating the attributes described in Section 2.2.1 for both the origin and destination hexagons. The encoder uses residual

[1]IRIS (**I**lots **R**egroupés pour l'Information **S**tatistique) is the basic sub-municipal statistical unit used by INSEE, designed to contain roughly 2,000 residents [INSEE, 2016].

blocks with dropout regularization to mitigate overfitting to territory-specific patterns and promote generalization across diverse territorial contexts.

The encoder operates in two stages (Fig. 1). During ***pre-training***, data from all source territories are pooled into a single training set, with each sample annotated with a domain label used by the adversarial classification head (Section 2.2.3) to encourage territory-agnostic representations. The encoder is thus jointly optimized across source territories, capturing generalizable patterns of DRT demand and travel time variability. During ***fine-tuning***, the pre-trained encoder is initialized with the learned parameters and updated on target territory data, which may be scarce. We evaluate the framework under a ***zero-shot*** setting (no target data) and multiple ***few-shot*** settings (varying amounts of target data).

During inference, the framework is applied to all possible OD pairs in the target territory. For each pair, the model predicts the trip activation probability and, for active pairs, the parameters μ and σ of the travel time distribution, which are subsequently used to compute accessibility indicators using Eq. 1.

2.2.3 Adversarial Domain Classification Head

The adversarial domain classification head encourages the encoder to produce domain-invariant feature vectors that are indistinguishable across source territories. This is achieved through a Gradient Reversal Layer (GRL) [Ganin et al., 2016], which during the backward pass negates and scales the gradients flowing from the domain classifier into the encoder by an annealing coefficient $\lambda \in [0,1]$, pushing the encoder to suppress territory-specific features while retaining information useful for downstream tasks. The head produces a probability distribution over the M source domains and is trained using the cross-entropy loss:

$$\mathscr{L}_{\text{domain}} = -\sum_{m=1}^{M} d_m \log \hat{d}_m \tag{2}$$

where $d_m \in \{0,1\}$ is the ground-truth domain indicator and $\hat{d}_m \in [0,1]$ is the predicted domain probability, with $\sum_{m=1}^{M} \hat{d}_m = 1$. This loss is minimized by the classifier but, due to gradient reversal, maximized from the encoder's perspective, driving domain confusion. This head is active exclusively during pre-training and is discarded during fine-tuning on the target territory.

2.2.4 OD Trip Activation Head

The OD trip activation head predicts whether a DRT trip is likely to occur between a given OD pair at a given departure time t. As noted earlier, this is formulated as a binary classification task: the head outputs a probability $\hat{p} \in [0,1]$ indicating the likelihood of trip activation, which is then thresholded to produce a binary prediction $\hat{y} \in \{0,1\}$, where 1 indicates an active OD pair and 0 indicates an inactive one. The head consists of a linear layer followed by a ReLU activation, a final linear layer, and a Sigmoid activation that maps the scalar output to $\hat{p} \in [0,1]$. Since the task is binary classification, the head is trained using the binary cross-entropy loss:

$$\mathscr{L}_{\text{activation}} = -[y \log \hat{p} + (1-y)\log(1-\hat{p})] \tag{3}$$

where $y \in \{0,1\}$ is the ground-truth activation label and $\hat{p} \in [0,1]$ is the predicted activation probability. The predicted activation serves a dual purpose within the framework: it acts as a proxy for DRT travel demand, and simultaneously gates the travel time estimation head, which is only applied to OD pairs predicted as active.

2.2.5 Travel Time Distribution Estimation Head

The travel time estimation head predicts the distribution of travel time for a given active OD pair. Rather than producing a point estimate, the head outputs two parameters: the mean $\hat{\mu}$ and standard deviation $\hat{\sigma}$, which are used to parameterize a Normal distribution $\mathscr{N}(\hat{\mu}, \hat{\sigma}^2)$ of the predicted travel time. Following Ourahou et al. [2025], the Normal distribution is adopted to model the stochastic nature of DRT travel times.

The head consists of a linear layer followed by a ReLU activation, a Dropout layer for regularization, and a final linear layer producing two scalar outputs. A Softplus activation is applied to $\hat{\sigma}$ to ensure strict positivity, while $\hat{\mu}$ is left unconstrained. The model is trained by minimizing the Negative Log-Likelihood (NLL), which jointly penalizes inaccurate mean predictions and poorly calibrated uncertainty estimates, encouraging the head to produce well-calibrated distributional outputs:

$$\mathscr{L}_{\text{travel time}} = -\log \mathscr{N}(T_{uv} \mid \hat{\mu}, \hat{\sigma}^2) = \frac{1}{2}\log(2\pi\hat{\sigma}^2) + \frac{(T_{uv}-\hat{\mu})^2}{2\hat{\sigma}^2} \tag{4}$$

where T_{uv} is the observed travel time for the OD pair (u,v), $\hat{\mu}$ is the predicted mean, and $\hat{\sigma} > 0$ is the predicted standard deviation. Note that this head is only applied to active OD pairs ($y = 1$), as travel time is undefined for inactive pairs.

2.3 Unified Modelling of DRT with PT

To integrate DRT into a multimodal accessibility framework, we adopt the Time-Expanded Graph (TEG) approach of Ourahou et al. [2025], which augments the PT TEG with virtual lines representing potential DRT trips, enabling centroid-to-centroid multimodal journeys combining walking, DRT, and PT in up to five trip patterns: Walking+DRT+Walking, Walking+PT+Walking, Walking+DRT+PT+Walking, Walking+PT+DRT+Walking, and Walking+DRT+PT+DRT+Walking.

In the original method, DRT virtual lines originate and terminate at barycentres of historical pick-up and drop-off locations. Since the target territory may have no historical DRT data, barycentres cannot be computed. We therefore substitute hexagon centroids for barycentres, connecting DRT virtual lines directly between hexagon centroids, which simplifies trip patterns while remaining consistent with the hexagonal spatial representation used throughout the framework.

The travel time of each DRT virtual line is set to the 95th percentile of the predicted Normal distribution $\mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$, providing a conservative estimate that accounts for travel time variability and prevents overestimation of accessibility. Only OD pairs predicted as active ($\hat{y} = 1$) are assigned DRT virtual lines; inactive pairs are excluded. The resulting TEG is converted into a GTFS feed and used to compute Eq. 1.

2.4 Accessibility Estimation Baseline and Evaluation Protocol

Since ground-truth DRT accessibility values are not directly observable — accessibility is a derived quantity depending on estimated travel times and demand — a reference baseline must be constructed for evaluation. To this end, we train a version of the DAMTL framework using only data from the target territory, without domain adaptation or transfer from source territories. In this configuration, the domain classification head is discarded, as there is a single territory and domain labels are not applicable. The resulting single-territory model is then used to perform inference on the same target territory.

The accessibility estimates produced by this single-territory model serve as a proxy for the reference accessibility, under the assumption that a model trained exclusively on local data provides a closer approximation to the true accessibility than a transferred model. This assumption holds when sufficient local data are available, and may not generalise to extremely data-scarce settings. While this reference is not the true ground truth, it represents the best achievable estimate given the available data, and constitutes a meaningful and internally consistent benchmark against which the zero-shot and few-shot transfer scenarios can be evaluated.

3 Results

3.1 Study Areas

The DAMTL framework is applied to three French metropolitan areas: Strasbourg and Lille as source territories, and Orléans as the target. In all three areas, the DRT service is operated by Padam Mobility [Padam Mobility, 2022] and requires pre-booking prior to departure.

Lille Metropolis (Fig. 2a) has a population of approximately 1,179,000 across 95 municipalities [INSEE, 2019]. Its DRT service, Ilévia Réservation, has been operational since 2019, *connecting peripheral municipalities to major PT corridors* [Ilévia, 2025], with 203,860 trips recorded between 2019 and 2025.

Strasbourg Metropolis (Fig. 2b) has a population of approximately 517,000 across 33 municipalities [INSEE, 2022]. Its DRT service, Flex'hop, has been operational since 2019, *operating as a zonal stop-to-stop service* across 25 municipalities [CTS, 2025], with 1,732,034 trips recorded between 2019 and 2025.

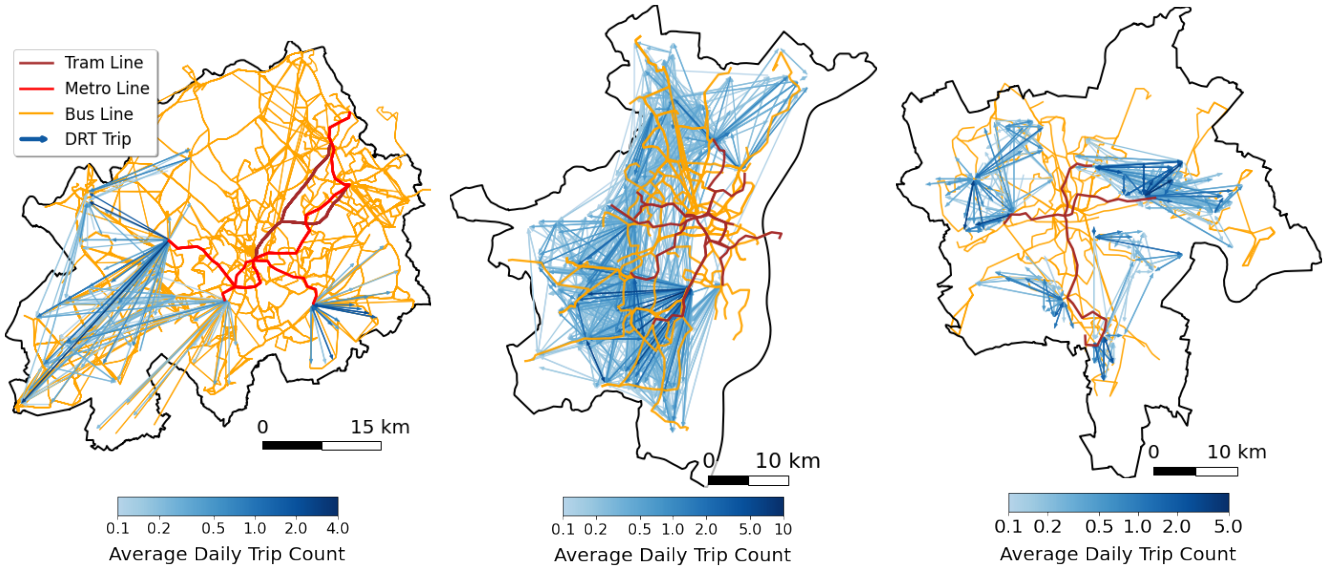
Orléans Metropolis (Fig. 2c), the target territory, has a population of approximately 293,000 across 22 municipalities [INSEE, 2022]. Its DRT service, operational since 2018, *combines first- and last-mile feeder trips with direct origin-to-destination journeys*, with 1,036,437 trips recorded between 2018 and 2025.

Fig. 2 shows the average daily number of DRT trips and their spatial distribution across the three territories, highlighting strong directional patterns either toward main PT corridors or between peripheral zones. The hourly demand profiles (Fig. 3) exhibit consistent temporal dynamics, with recurring peaks around 7 a.m., 12 p.m., and 4–5 p.m., although demand intensity varies across cities, with Strasbourg showing the highest overall demand and Lille the lowest.

3.2 OD Activation Classification Performance

The framework is evaluated through a set of scenarios reflecting the full transfer learning pipeline. For both source and target territories, data are split chronologically, with the first 80% used for training and the remaining 20% (most recent data) used for validation. For the target territory, few-shot training sets are constructed from the most recent data, ranging from 1 day to 25 weeks (e.g., the 1-week shot corresponds to the latest week of data). The training configurations are reported in Table 1.

Fig. 4 reports the ROC curves and corresponding AUC scores across all evaluated transfer scenarios. The reference baseline, trained exclusively on Orléans Metropolis data, provides a strong upper bound for the transfer experiments. Under zero-shot transfer, the model performs close to random, confirming a significant distribution shift between source and target territories. Few-shot fine-tuning progressively addresses this shift, with classification performance improving steadily as the share of



(a) Lille Metropolis **(b)** Strasbourg Metropolis **(c)** Orléans Metropolis

Figure 2. Spatial distribution of PT networks and average daily number of DRT trips.

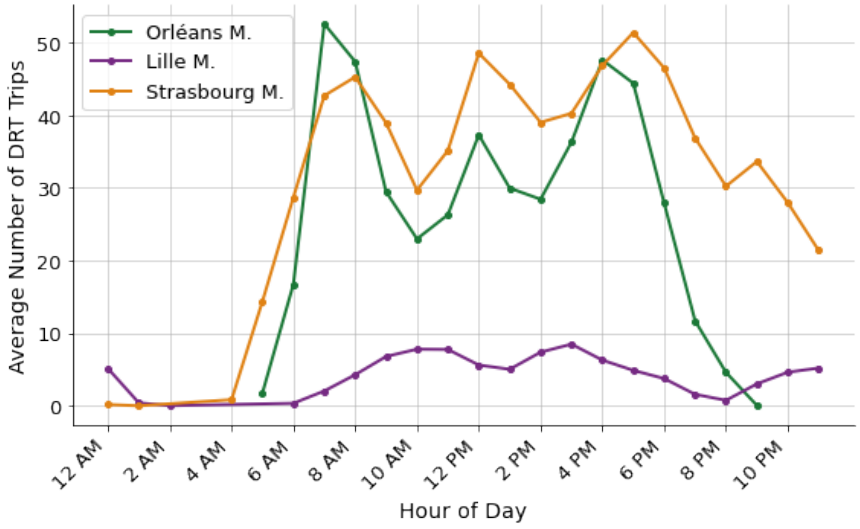


Figure 3. Average hourly DRT trip counts per metropolitan area.

target training data increases. **Notably, the 5 weeks-shot regime already approaches the reference baseline**, suggesting that a moderate amount of local data is sufficient to recover most of the classification performance. At 12 weeks-shot and 25 weeks-shot, the model surpasses the reference baseline, which is expected: the transferred model benefits from both the rich feature representations learned across source territories during pre-training and the local fine-tuning signal, whereas the reference baseline relies solely on local data without any cross-territory knowledge transfer.

Table 1. Summary of experimental parameters.

Category	Parameter	Value
Accessibility computation	Hexagon side length	2 km
	Min. population density	6 inhab./km²
	Maximum travel time τ (Eq. 1)	30 min
	Travel time percentile	95th
Dataset construction	Negative-to-positive DRT trip ratio	10:1
	DRT trips – Strasbourg M. (after 10:1)	1,732,034 × 11
	DRT trips – Lille M. (after 10:1)	203,860 × 11
	DRT trips – Orléans M. (after 10:1)	1,036,437 × 11
Model architecture	Input → Residual block 1	input → 256
	Residual block 1 → Residual block 2	256 → 128
	Activation head	128 → 1
	Travel time head	128 → 64 → 2
	Domain head	128 → 64 → 1
	Dropout rate	0.1
	GRL λ	Annealed from 0 to 1 during pre-training
Training	Optimizer	Adam
	Learning rate	1e-4
	Batch size	1024
	Pre-training epochs	100
	Fine-tuning epochs	20
Transfer regimes	zero-shot	0 trips
	1-day-shot	1454 trips
	1-week-shot	10,178 trips
	5-weeks-shot	50,890 trips
	12-weeks-shot	122,136 trips
	25-weeks-shot	254,450 trips
Data split	Training set	80%
	Validation set	20%

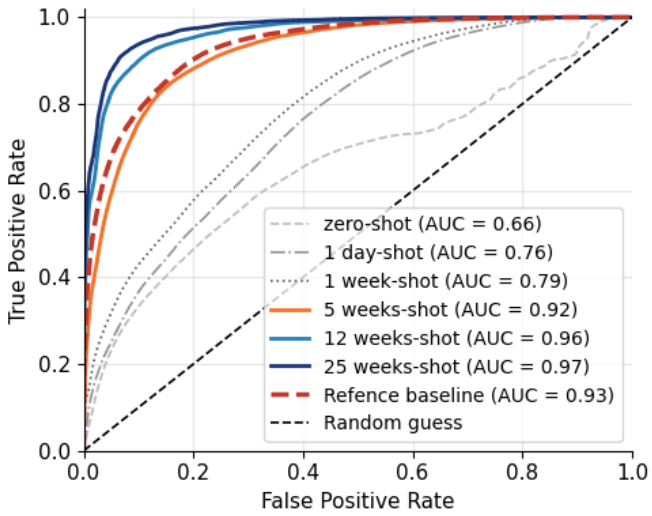


Figure 4. Receiver Operating Characteristic (ROC) curves and corresponding Area Under the Curve (AUC) scores of the OD activation classifier across transfer learning scenarios: zero-shot, few-shot fine-tuning, and reference baseline on Orléans Metropolis.

3.3 Accessibility Estimation Results

Accessibility indicators are computed by enriching the PT TEG with virtual walking connections and DRT virtual lines, as explained in Section 2.3. Accessibility is evaluated at a departure time of 7 a.m., reflecting peak morning conditions. Accessibility scores are also computed at 10 a.m. and 4 p.m.; however, for brevity, this section focuses on the 7 a.m. results. The travel time threshold $\tau = 30$ minutes is chosen to align with typical commuting times in the three territories. To evaluate the quality of the transferred accessibility estimates, we compute the absolute difference in the percentage of jobs accessible between the predicted and reference accessibility scores for each hexagon:

$$\delta = \mathrm{Acc}_{\mathrm{pred}}(\%) - \mathrm{Acc}_{\mathrm{ref}}(\%) \tag{5}$$

where $\mathrm{Acc}_{\mathrm{pred}}(\%)$ and $\mathrm{Acc}_{\mathrm{ref}}(\%)$ are the predicted and reference job accessibility scores expressed as a percentage of total jobs accessible within the travel time threshold τ, respectively.

The spatial distribution of differences is illustrated in Figs. 5 and 6 for different fine-tuning scenarios. In the zero-shot setting, large differences are concentrated in the central urban area, particularly around Orléans and Fleury-les-Aubrais, which correspond to zones with the highest job density, and are therefore most sensitive to errors in the transferred demand and travel time estimates. As the fine-tuning fraction increases, these differences progressively diminish. From the 5 weeks-shot regime onwards, the spatial pattern of differences becomes notably sparser and less intense, with only a few isolated hexagons retaining non-negligible deviations. At higher fine-tuning fractions, the predicted accessibility maps closely resemble the reference baseline across the entire territory, confirming that even a moderate amount of local data is sufficient to spatially align the transferred estimates with the reference.

Figure 7 shows the mean absolute difference across hexagons as a function of the fine-tuning data fraction. In the zero-shot setting, the mean error is 5.82 percentage points (i.e., jobs), reflecting the difficulty of transferring accessibility estimates without any target territory data. As the fine-tuning fraction increases, the error decreases substantially, reaching 1.21 percentage points at 25 weeks-shot.

To quantify the efficiency of each fine-tuning regime, we define a recovery rate ρ that measures the share of total recoverable improvement achieved at each fine-tuning level:

$$accessibility\rho = \frac{E_{\text{zero-shot}} - E_{\text{current}}}{E_{\text{zero-shot}} - E_{25-weeks-shot}} \times 100 \tag{6}$$

where E denotes the mean absolute difference in percentage points of accessible jobs, $E_{\text{zero-shot}}$ and $E_{\text{25-weeks-shot}}$ are its values at zero-shot and 25 weeks-shot respectively, and E_{current} is its value at the fine-tuning regime under evaluation.

The improvement exhibits clear diminishing returns as data availability increases. **The 5 weeks-shot regime alone recovers 70.1% of the total possible improvement, while the entire transition from 5 weeks-shot to 25 weeks-shot yields only an additional 29.9% of improvement.** The IQR also narrows with increasing data availability, indicating not only lower average errors but also more spatially consistent estimates across hexagons.

These results demonstrate that leveraging data from other territories and transferring the learned knowledge to a new territory can yield accessibility estimates close to the reference baseline, even when only a very limited amount of local operational data is available.

4 Discussion and Limitations

Estimating accessibility in territories where DRT services are newly deployed or operate with limited historical data remains an open methodological challenge. Existing approaches either rely on fully observed demand and schedules or resort to simplified assumptions that fail to capture the stochastic nature of DRT operations. The proposed DAMTL framework addresses this gap by jointly estimating OD activation and travel time distributions through adversarial domain adaptation, enabling accessibility computation in target territories with little to no local data.

Several limitations should be acknowledged. First, the reference accessibility is itself a model-based estimate rather than observable ground truth, meaning systematic bias in the reference propagates into the evaluation metric. Second, the framework uses travel time as the sole generalised cost, omitting fares, transfer penalties, and discomfort; while DRT and PT fares are aligned in the study areas, future work should incorporate multi-dimensional cost measures. Third, the pre-booking operational model studied yields systematically lower waiting times than fully on-demand systems, potentially producing optimistic accessibility estimates; incorporating dynamic waiting-time distributions would broaden applicability. Fourth, PT travel time variability is not modelled, which could affect reliability of estimates in territories with irregular service. Fifth, OD activation is treated as binary rather than volumetric, so DRT virtual lines are added for all predicted active pairs regardless of fleet constraints, potentially overestimating accessible destinations. The framework is modular and indicator-agnostic: the cumulative job accessibility measure adopted here can be replaced by gravity-based, opportunity variety, or alternative

threshold-based indicators without modifying the estimation pipeline. Future work will prioritise sensitivity analysis over key parameters (Travel time threshold, hexagon resolution, number of source territories) and validation across additional target territories with diverse urban configurations and DRT operational models.

5 Conclusion

This paper proposes the DAMTL framework for estimating DRT accessibility in data-scarce territories through transfer learning. Applied to Orléans Metropolis using Lille and Strasbourg as sources, results show that zero-shot transfer already produces coherent accessibility patterns, and that as little as 5 weeks of local DRT operation recover 70.1% of the total recoverable improvement. These findings demonstrate that transport authorities can leverage operational data from existing DRT deployments to anticipate accessibility performance before a sufficient local dataset becomes available.

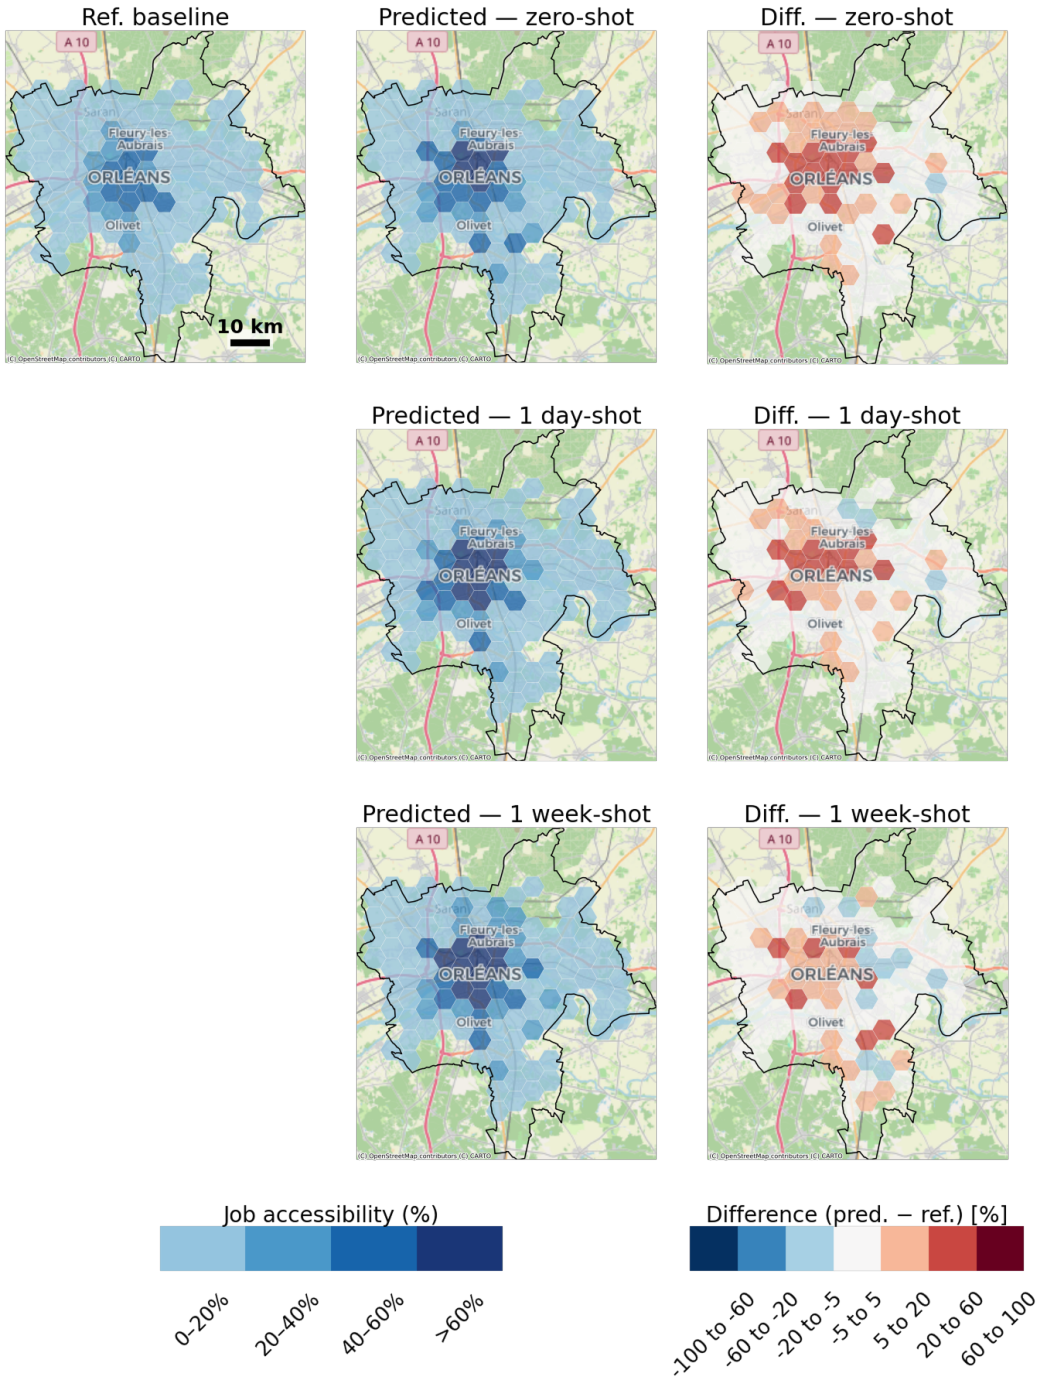


Figure 5. Comparison of estimated and reference baseline accessibility scores across fine-tuning regimes for Orléans Metropolis at 7 a.m. (Part 1).

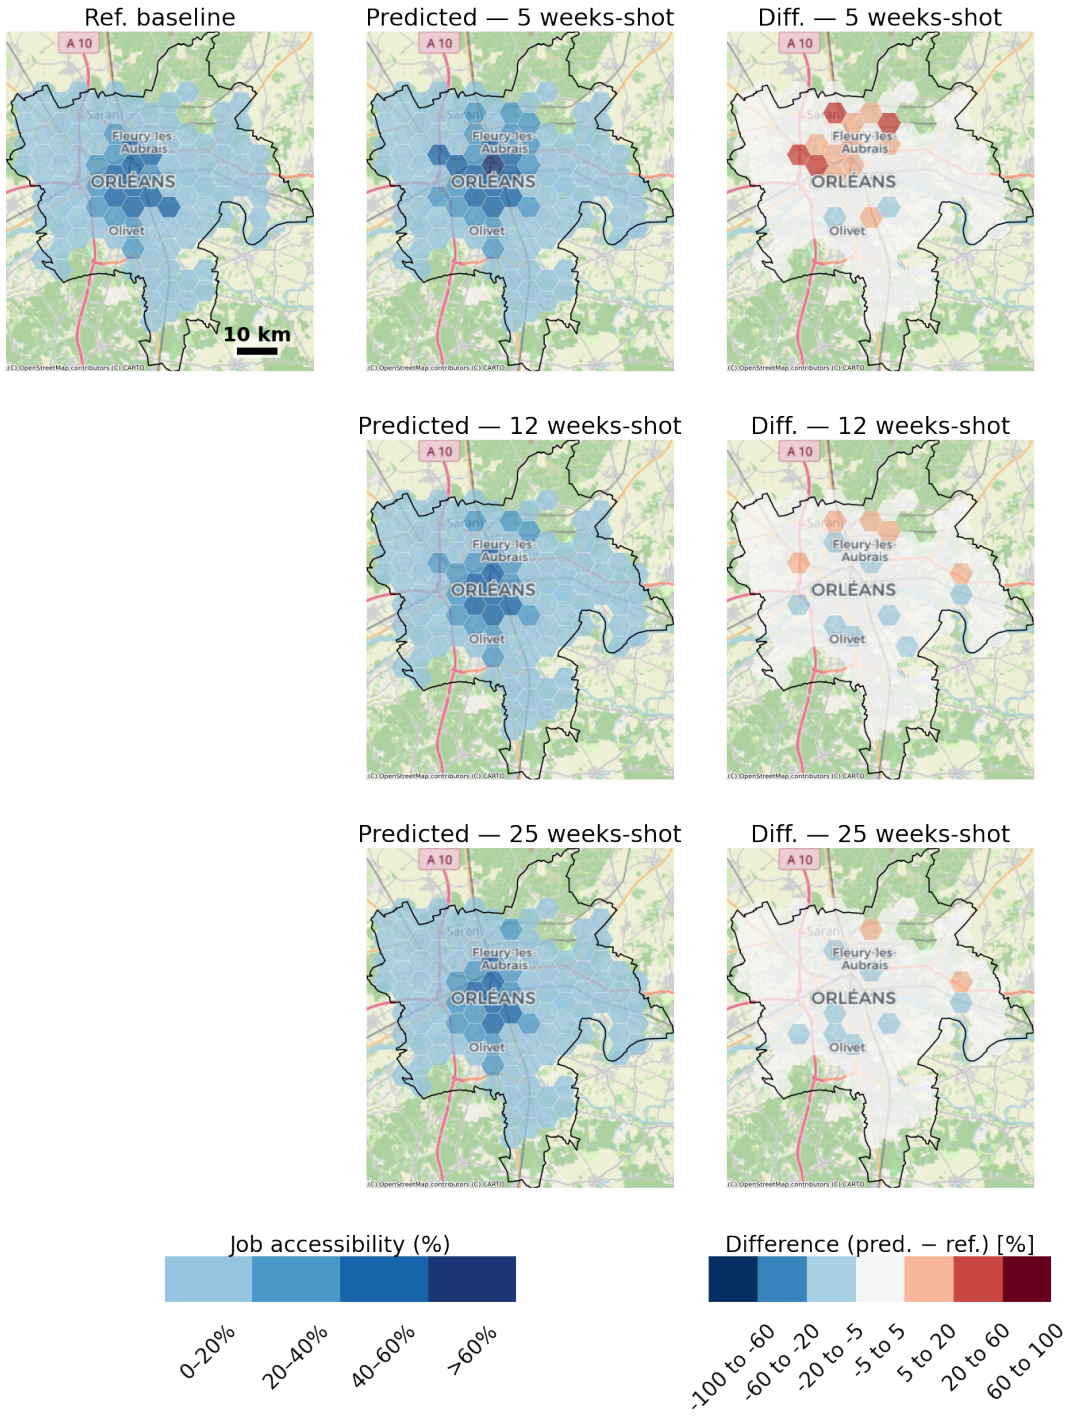


Figure 6. Comparison of estimated and reference baseline accessibility scores across fine-tuning regimes for Orléans Metropolis at 7 a.m. (Part 2).

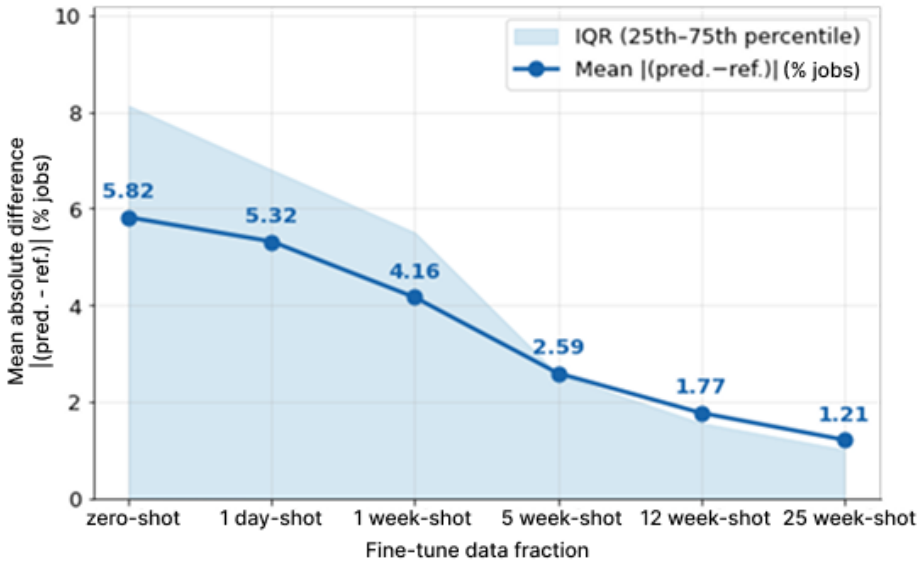


Figure 7. Mean absolute difference and interquartile range of accessibility estimates as a function of the fine-tuning data fraction.

References

M. J. Alonso-González, T. Liu, O. Cats, N. van Oort, and S. Hoogendoorn. The potential of demand-responsive transport as a complement to public transport: An assessment framework and an empirical evaluation. *Transportation Research Record*, 2672(8):879–889, 2018. doi: 10.1177/0361198118790842. URL https://doi.org/10.1177/0361198118790842.

I. Biazzo, B. Monechi, and V. Loreto. General scores for accessibility and inequality measures in urban areas. *Royal Society Open Science*, 6(10):191014, 2019.

K. Bimpou and N. S. Ferguson. Dynamic accessibility: Incorporating day-to-day travel time reliability into accessibility measurement. *Journal of Transport Geography*, 96:103170, 2021.

C. K. V. Braga, C. F. G. Loureiro, and R. H. M. Pereira. Evaluating the impact of public transport travel time inaccuracy and variability on socio-spatial inequalities in accessibility. *Journal of Transport Geography*, 109:103590, 2023.

G. Calabrò, A. Araldo, M. Ben-Akiva, et al. Adaptive transit design: Optimizing fixed and demand responsive multi-modal transport via continuous approximation. *Transportation Research Part A: Policy and Practice*, 2023.

M. W. Conway, A. Byrd, and M. Van Eggermond. Accounting for uncertainty and variation in accessibility metrics for public transport sketch planning. *Journal of Transport and Land Use*, 11(1):541–558, 2018.

CTS. Rapport d'activité et chiffres clés 2024-2025. Technical report, CTS, Strasbourg, France, 2025. URL https://www.cts-strasbourg.eu/en/know-us/the-essentials-of-the-cts/who-are-we-/our-key-figures/.

G. Currie and N. Fournier. Why most drt/micro-transits fail–what the survivors tell us about progress. *Research in Transportation Economics*, 83:100895, 2020.

S. Diepolder, A. Araldo, T. Chouaki, S. Maiti, S. Hörl, et al. On the computation of accessibility provided by shared mobility. In *11th Symposium of the European Association for Research in Transportation (hEART)*, Zurich, Switzerland, 2023. URL https://hal.archives-ouvertes.fr/hal-04154712v2.

Severin Diepolder, Andrea Araldo, Tarek Chouaki, Santa Maiti, Sebastian Hörl, and Constantinos Antoniou. Quantifying the improvement of accessibility achieved via shared mobility on demand. *Transportation*, 2025. doi: 10.1007/s11116-025-10662-1.

S. Farber, M. Z. Morang, and M. J. Widener. Temporal variability in transit-based accessibility to supermarkets. *Applied Geography*, 53:149–159, 2014.

A. Fielbaum and J. Alonso-Mora. Unreliability in ridesharing systems: Measuring changes in users' times due to new requests. *Transportation Research Part C: Emerging Technologies*, 121:102831, 2020.

K. Fransen, T. Neutens, S. Farber, P. De Maeyer, G. Deruyter, and F. Witlox. Identifying public transport gaps using time-dependent accessibility levels. *Journal of Transport Geography*, 80:102501, 2019.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):59–98, 2016.

Walter G. Hansen. How accessibility shapes land use. *Journal of the American Institute of Planners*, 25(2):73–76, 1959.

Mingzhuang Hua, Francisco Camara Pereira, Yu Jiang, Xuewu Chen, and Junyi Chen. Transfer learning for cross-modal demand prediction of bike-share and public transit. *Journal of Intelligent Transportation Systems*, 28(5):640–653, 2024. doi: 10.1080/15472450.2024.2371913.

Ilévia. Le réseau ilévia, 2025. URL https://www.ilevia.fr/fr/decouvrir-le-reseau/90/Le-reseau-ilevia. Accessed 5 August 2025.

INSEE. Définition - iris. https://www.insee.fr/fr/metadonnees/definition/c1523, 2016. Accessed: 03 Nov. 2025.

INSEE. La métropole européenne de lille : un quart des emplois de la région. Technical report, French National Institute of Statistics and Economic Studies, 2019. URL https://www.insee.fr/fr/statistiques/6794926.

INSEE. Comparateur de territoires : Eurométropole de Strasbourg (246700488) et Orléans Métropole (200042268), 2022. URL https://www.insee.fr/fr/statistiques/1405599?geo=EPCI-246700488+FE-1.

INSEE. Insee - national institute of statistics and economic studies. https://www.insee.fr/fr/accueil, 2025. Accessed 3 September 2025.

S. Jiang, W. Guan, Z. He, and L. Yang. Measuring taxi accessibility using grid-based method with trajectory data. *Sustainability*, 10(9):3187, 2018.

Jintao Ke, Hongyu Zheng, Hai Yang, and Xiqun Michael Chen. Short-term forecasting of passenger demand under on-demand ride services: A spatio-temporal deep learning approach. *Transportation Research Part C: Emerging Technologies*, 85:591–608, 2017. doi: 10.1016/j.trc.2017.10.016.

S. Li, H. Duan, T. E. Smith, and H. Hu. Time-varying accessibility to senior centers by public transit in philadelphia. *Journal*

of Transport Geography, 82:102566, 2020.

Karen Lucas. Transport and social exclusion: Where are we now? *Transport Policy*, 20:105–113, 2012. doi: 10.1016/j.tranpol.2012.01.013. URL https://doi.org/10.1016/j.tranpol.2012.01.013.

E. Miller. Measuring accessibility. Discussion paper, International Transport Forum, Organisation for Economic Co-operation and Development (OECD), 2020.

Mohamed Ourahou, Andrea Araldo, Louis Zigrand, Giovanni Calabrò, and Aline Carneiro Viana. Measuring accessibility in demand-responsive transport: A statistical method capturing travel time reliability. SSRN, 2025. URL https://ssrn.com/abstract=5465463. Available at SSRN.

Padam Mobility. Use case, 2022. URL https://www.padam-mobility.com/.

Tao Peng, Jie Gao, and Oded Cats. Uncertainty-aware probabilistic travel demand prediction for mobility-on-demand services. *Transportation Research Part C: Emerging Technologies*, 181:105383, 2025. doi: 10.1016/j.trc.2024.105383.

J. Pucher. Public transportation. In S. Hanson and G. Giuliano, editors, *The Geography of Urban Transportation, 3rd edition*. Guilford Press, New York, 2004.

Sudipta Roy, Bat-hen Nahmias-Biran, and Samiul Hasan. Spatial transferability of machine learning based models for ride-hailing demand prediction. *Transportation Research Part A: Policy and Practice*, 193:104413, 2025. doi: 10.1016/j.tra.2025.104413.

P. Saeidizand, K. Fransen, and K. Boussauw. Revisiting car dependency: A worldwide analysis of car travel in global metropolitan areas. *Cities*, 2022.

K. Viergutz and C. Schmidt. Demand responsive - vs. conventional public transportation: A matsim study about the rural town of colditz, germany. *Procedia Computer Science*, pages 69–76, 2019.