

DeepMEND: Reliable and Scalable Network Metadata Geolocation from Base Station Positions

Orlando E. Martínez-Durive^{*†}, Stefanos Bakirtzis[‡], Cezary Ziemlicki[§],
Jie Zhang[¶], Ian James Wassell[‡], Marco Fiore^{*}

^{*}IMDEA Networks Institute, Spain, [†]Universidad Carlos III de Madrid, Spain, [‡]University of Cambridge, United Kingdom

[§]SENSE / Orange Innovation, France [¶]University of Sheffield, Dept. of Electronic and Electrical Engineering, United Kingdom
{orlando.martinez, marco.fiore}@imdea.org, {ssb45, ijw24}@cam.ac.uk, cezary.ziemlicki@orange.com, jie.zhang@sheffield.ac.uk

Abstract—Metadata geolocation, *i.e.*, mapping information collected at a cellular Base Station (BS) to the geographical area it covers, is a central operation in the production of statistics from mobile network measurements. This task requires modeling the probability that a device attached to a BS is at a specific location, and is presently addressed with simplistic approximations based on Voronoi tessellations. As we show, Voronoi cells exhibit poor accuracy compared to real-world geolocation data, which can, in turn, reduce the reliability of research results. We propose a new approach for data-driven metadata geolocation based on a teacher-student paradigm that combines probabilistic inference and deep learning. Our DEEPMEND model: (i) only needs BS positions as input, exactly like Voronoi tessellations; (ii) produces geolocation maps that are 56% and 33% more accurate than legacy Voronoi and their state-of-the-art VoronoiBoost calibration, respectively; and, (iii) generates geolocation data for thousands of BSs in minutes. We assess its accuracy against real-world multi-city geolocation data of 5,947 BSs provided by a network operator, and demonstrate the impact of its enhanced metadata geolocation on two applications use cases.

I. INTRODUCTION

The metadata generated by mobile networks comprises a variety of Key Performance Indicators (KPIs, *e.g.*, call detail records, data traffic volumes, mobile service demands, radio channel statistics, or user presence, among many others) that provide rich information about the movement, communication, activities, and interests of large populations of subscribers, with high spatiotemporal resolution and at broad geographical scales. The data has applications in many domains, including networking, demography, geography, sociology, or epidemiology, where it enables dependable quantitative analyses of large-scale phenomena. Examples include: the study of time-varying population densities [1], [2], land use [3], or transportation system planning [4]; the mapping of urban transformations [5], pollution [6], digital divides [7], poverty [8] and social inequality [9]; the dynamics of natural disasters [10] or infectious diseases [11], and the effectiveness of their containment policies [12]; the characterization of mobile traffic [13], [14] and application usage [15], and the prediction of their spatiotemporal fluctuations [16], just to cite a few. The value of network metadata is driving the creation of dedicated statistics production processes [17] and open-source toolkits [18].

Mobile network metadata is today geo-referenced at the highest resolution available in production systems, *i.e.*, at the level of Base Stations (BSs). Hence, a metadata record generated by one device is associated to the geographical site of the BS serving the device when the record is produced.




Fig. 1: Metadata geolocation of a real-world BS, obtained from (a) processing of accurate propagation simulations and field measurements by a network operator, (b) a legacy Voronoi tessellation of space, (c) our proposed DEEPMEND.

Yet, any network metadata-driven analysis must be carried out over the continuous geographical space, and not on the discrete BS locations. The mapping of BS-referenced data to the continuous space is then an essential step in any mobile network metadata processing pipeline [17]. The task is called *metadata geolocation*, and consists of *modeling the probability that the metadata attributed to one BS has been generated by a device at a specific geographical location served by that BS*.

Accurate metadata geolocation can be derived in current 4G/5G deployments by mobile network operators via a combination of extensive client-side measurements and radio propagation models informed by complete Radio Access Network (RAN) configurations, as we expound in Section III-A. An example of geolocation derived in this way for one real-world BS of Orange, is shown in Figure 1a. The target BS is positioned at the center of the plot, and colors denote different probabilities that the metadata recorded at the BS is generated at each geographical location. While highly accurate, such an approach to metadata geolocation requires access to confidential information about the network infrastructure as well as costly and time-consuming measurement campaigns. Ultimately, the approach can only be run by mobile operators, and the resulting maps are typically kept confidential.

The unavailability of precise geolocation forces the research community to resort to simple approximations to geolocate the network metadata. Approaches based on Voronoi tessellations are the common practice, *e.g.*, they are adopted by all works and tools referenced earlier [1]–[18]. The strategy uniformly spreads the metadata associated to one BS over the surface of its Voronoi cell, *i.e.*, the locus of spatial points that are closer to the target BS than to any other BSs: an example of geolocation based on a legacy Voronoi tessellation is illustrated in Figure 1b for the same BS of Figure 1a. Voronoi

tessellations are simple to compute and, importantly, only require the geographical coordinates of the BSs as input. Yet, the limitations of a Voronoi-based model are apparent even in the example of Figure 1b, and has been criticized [19], [20].

The aim of this paper is *democratizing dependable metadata geolocation*, by proposing DEEPMEND, a model that combines the minimal input requirement of a Voronoi tessellation (*i.e.*, BS locations only) with the high accuracy of a complex model based on proprietary RAN information. Our solution performs deep metadata geolocation with knowledge distillation: a sample of the geolocation produced by our tool is illustrated in Figure 1c, for the same BS of Figure 1a.

We publicly release¹ our ready-to-use implementation of DEEPMEND, so that it can benefit the community and possibly be integrated into emerging tools for producing official statistics from network metadata [17], [18].

II. RELATED WORK

Next, we review existing approaches to metadata geolocation, which we group according to their methodology.

Client-side measurements. Precise mappings of the real-world metadata geolocation of BSs inherently need information collected at the User Equipment (UE). While complete UE-side measurements may be feasible in circumstantial scenarios (*e.g.*, for a few BS), they do not scale to large radio access infrastructures covering cities or countries, which are typical settings for network metadata-driven analyses [21]. A way to alleviate the problem is collecting UE-side geolocation samples at a limited subset of locations, *e.g.*, via wardriving, and then employing signal reconstruction techniques to retrieve the complete maps [22], [23]. Yet, the approach still requires massive campaigns to gather enough samples over vast multi-city regions like those we target. As a result, no scientific study based on mobile network metadata has employed client-side measurement campaigns for geolocation.

Propagation-based modeling. A more scalable way to obtain accurate metadata geolocation maps is to rely on radio-frequency (RF) signal propagation models to compute the strength of the received signal from a BS at each point in space and then post-process such information to derive the actual geolocation maps. However, the approach necessitates (*i*) complete data about the wireless propagation environment, (*ii*) access to reliable RF propagation solvers to estimate the wireless channel characteristics, and (*iii*) full information about the RAN configuration that includes confidential parameters, such as the modulation schemes and code rates used by each BS. Even recent advances in deep learning that have lowered the computational barrier to RF propagation modeling [24], [25] still need data about the terrain, building layout, operating frequency, transmitted power, modulation and coding schemes, antenna height, azimuth, tilt or radiation pattern. These barriers, jointly to the strong domain knowledge needed to use complex RF propagation solvers, determine that propagation-based methods are never used in academic research on knowledge inference from network metadata.

Voronoi tessellations. The complexity of the approaches above have paved the way for the wide adoption of a much simpler strategy for metadata geolocation based on a Voronoi tessellation of space. Straightforward computation and minimum input data limited to the BS locations make Voronoi cells the de-facto standard for geolocation in the literature.

Variations to the baseline Voronoi model have also been proposed, by: (*i*) shifting the Voronoi anchors along the BS azimuth or using knowledge of the BS coverage area boundaries to improve the anchor positioning [26]–[28]; (*ii*) extending Voronoi calculations with auxiliary data sources, such as land use, to improve the geographic location of devices [29]; or, (*iii*) embedding technology parameters for each BS, such as sectoring, transmission power, or path loss information to weigh the legacy Euclidean distance measures in a multiplicatively weighted Voronoi tessellation [30]. Yet, all these models had limited success, as they add significant computational complexity but do not overcome the main limitations of the original Voronoi tessellation: *e.g.*, they still assume a uniform probability of geolocation in space (see Figure 1b), which does not hold in practice (see Figure 1a).

The state-of-the-art Voronoi variant is VoronoiBoost, which scales the baseline Voronoi cell of a BS so that it better matches the area containing a target fraction of the geolocation probability [20]. While it is still relatively inaccurate, as discussed in Section III, VoronoiBoost is the closest work to ours, and we provide a complete comparative evaluation against DEEPMEND in Section V.

Probabilistic density models. An alternative strategy to Voronoi tessellation that has been proposed in the literature relies on a probability density inference (PDI) model [31]. The approach generates a statistical representation of the location of a UE given its distance from the BS of attachment, computing an approximate probabilistic metadata geolocation of the BS. The original PDI needs fine-tuning with measurement data collected in the target network, about the BSs transmitting power, the signal-to-noise and interference ratio (SINR) thresholds, interference from non-neighboring base stations, or round trip time (RTT): as these are typically not available to academic researchers, the method suffers from a similar accessibility barrier as proposals based on RF propagation solvers, and has thus met negligible adoption in the literature.

A simplified version of PDI that only requires BS locations has also been proposed [32]. This more essential model results in a Voronoi cell smoothing controlled by a power attenuation parameter, which requires again extensive field measurements to be properly configured. We prove in Section III that using the roughly calibrated values of the attenuation recommended for the model [31] for different urbanization levels result in poor geolocation approximations, and then also include this solution in the comparative evaluation of Section V.

Contribution. With respect to the literature above, we seek an original modeling approach that: (*i*) avoids the obstacles of the more accurate approaches, and does not rely on costly software packages, long and computationally expensive simulations or confidential information about the RAN infrastructure,

¹<https://github.com/nds-group/DeepMEND>

which currently curb usability by the research community; and, (ii) guarantees a representation of metadata geolocation that is as realistic as possible.

III. DATASET AND MOTIVATION

We first present the metadata geolocation data computed and made available by a network operator for our study, in Section III-A. We then introduce a similarity metric to compare metadata geolocation maps, in Section III-B. Finally, we quantitatively assess the reliability of current models for metadata geolocation based on BS locations, in Section III-C.

A. Ground-truth metadata geolocation data

Our work builds upon metadata geolocation maps for 5,947 outdoor BSs located in tens of urban, suburban, and rural areas of France, provided by Orange, a leading mobile network operator. The metadata geolocation dataset is derived from realistic spatial distributions of the *association probability* $P(i|\ell)$: these describe the probability that a UE located at position ℓ is attached to the i -th BS², hence $\sum_i P(i|\ell) = 1$. For each BS, the probability of association are provided over a spatial grid centered at the BS and composed of 600×600 location elements ℓ , each covering 0.01 km^2 .

The operator determines association probabilities through a three-step process. First, a proprietary propagation solver is utilized to estimate the received signal strength from each base station (BS) i at every location ℓ with high accuracy. Second, internally engineered priority tables are used to infer $P(i|\ell)$ based on the signal strength from all BSs covering location ℓ . Finally, extensive wardriving campaigns are conducted to collect client-side measurements, which are used to experimentally validate the association probabilities $P(i|\ell)$ and to potentially fine-tune the priority table weights.

The metadata geolocation $P(\ell|i)$, *i.e.*, the probability that an UE associated to BS i is located at ℓ , is then derived from the association probability $P(i|\ell)$ via Bayesian inference, as

$$P(\ell|i) = \frac{P(i|\ell)P(\ell)}{P(i)} = \frac{P(i|\ell)P(\ell)}{\sum_{\ell'} P(i|\ell')P(\ell')}, \quad (1)$$

where $P(\ell)$ is the prior probability that a UE is found at a specific grid position ℓ . As $P(\ell)$ is a time-varying distribution that depends on the user population mobility dynamics and is hard to determine a priori, we adopt the common assumption that the user density is uniform over the relatively limited area covered by overlapping BSs [33]. In this case, (1) reduces to $P(\ell|i) = P(i|\ell) / \sum_{\ell'} P(i|\ell')$, and $\sum_{\ell} P(\ell|i) = 1$.

B. Measuring metadata geolocation errors

A suitable error metric is required to compare the realistic metadata geolocation data described in Section III-A against approximations obtained by different modeling approaches. Given that geolocation is inherently stochastic, and the ground-truth $P(\ell|i)$ is a probability distribution, we adopt a metric

²We denote by i a whole BS, hence $P(i|\ell)$ is the total probability that UEs at ℓ associate to any of the carriers at the i -th BS site. We model geolocation at BS level rather than carrier level, as the former is much more popular in the literature. Yet, our model is easily generalized to carrier-level geolocation.




Fig. 2: Distributions of KS2 between the ground-truth metadata geolocation of each BS and Voronoi (blue), sPDI (orange), and VoronoiBoost (olive) models across urbanization levels.

from the field of probability. The following considerations drive our choice: (i) locations ℓ are placed in a discrete bi-dimensional space, making geolocation a multi-variate probability distribution; (ii) the metric must be defined for values where the reference probability is zero; and, (iii) the large support of 360,000 spatial elements in the geolocation map of each BS call for high computational efficiency.

The Kullback–Leibler and Jensen–Shannon divergences are not defined over zero values, while other common choices like the Wasserstein distance are computationally not viable. We select instead a generalization of the Kolmogorov–Smirnov test that is natively designed for bi-dimensional Cumulative Distribution Functions (CDFs) [34], and adopt an efficient version that meets our scaling needs [35]. As we do not seek to accept or reject a null hypothesis, we only employ the statistic returned by the bi-dimensional Kolmogorov–Smirnov test, *i.e.*, the maximum distance between the two-dimensional CDFs over all possible ordering, and not its p-value. The statistic, which we refer to as KS2 in the remainder of the document, is readily interpreted as an error metric: it ranges in $[0, 1]$, where 0 indicates that the metadata geolocation maps are identical, whereas 1 implies completely different representations.

C. Reliability of current models

We use the KS2 metric presented in Section III-B to perform a quantitative evaluation of the reliability of existing models for metadata geolocation. Adhering to the target of our study, we focus on models that are usable by the academic research community as they only need the locations of BSs as input. We thus evaluate (i) the legacy Voronoi tessellation, (ii) the trained implementation of the VoronoiBoost model [20], and (iii) the version of PDI that only uses BS location [32]. The latter model is denoted as sPDI in the following, and is configured by setting its single attenuation parameter to 5, 4, and 3 for BSs in urban, suburban, and rural locations, respectively, as indicated in the paper originally introducing the PDI model [31]; to this end, we determine the urbanization level of the area where each BS is positioned based on official demographic data [36].

Figure 2 summarizes the KS2 similarity measured between the ground-truth geolocation presented in Section III-A and the approximations returned by the three models above, for each BS in our dataset. The distribution of KS2 values is separated by urbanization levels, which yield a different parametrization of sPDI as well as heterogeneous accuracy results.

VoronoiBoost offers slightly lower KS2 errors, hence better accuracy than sPDI, which outperforms the baseline Voronoi tessellation. However, all these approximations still leave large margins for improvement: performance is especially poor with suburban and rural BSs, where the average error has high KS2 values above 0.5; for urban BSs, KS2 values are spread across the whole abscissa, implying that the geolocation output can be anywhere from reliable to entirely wrong.

IV. DEEPMEND

The previous analysis shows that available models generate unreliable mappings of metadata to the geographical space, which can undermine the credibility of downstream analyses. We propose DEEPMEND, a novel solution that follows a *knowledge distillation* (KD) approach [37], *i.e.*, employs a complex and computationally expensive model –the *teacher*– to derive soft labels that are then used to train a simpler model –the *student*. Both teacher and student in DEEPMEND are original models that are carefully tailored to solve the specific problem of generating realistic geolocation maps from BS locations, making our KD implementation unique.

The design of DEEPMEND is illustrated in Figure 3. At a high level, the teacher processes (i) BS locations and (ii) the ground-truth geolocation data presented in Section III-A, and distills so-called *soft labels*: these are approximate geolocation maps similar to the ground truth but less information-dense and easier to model. Then, the student uses these soft labels to learn to transform the BS locations into metadata geolocation maps that mimic the soft labels. It is worth noting that:

- *the teacher cannot operate as a generator of geolocation data from BS positions only since it requires ground-truth information as input*, hence its sole purpose is to support the learning process of the student during training;
- *during inference, the trained student can independently generate geolocation maps from BS coordinates only*, thus achieving our objective of a model that relies on the same input of legacy Voronoi-based approaches.

The rationale for adopting KD in the design of DEEPMEND is that the weak correlations between the tangled ground-truth metadata geolocation and the modest input represented by BS locations are too complex to learn in a direct way. By introducing a teacher model, we generate soft labels that are close to the target ground truth, yet more explicitly correlated to the BS positioning data, thus simplifying the learning process. To demonstrate the advantage of knowledge distillation and corroborate our design choice, in Section V-A, we will also compare DEEPMEND against a student-only solution that tries to learn to synthesize geolocation directly from the ground truth without the support of a teacher model.

A. Teacher model

Our teacher is a probabilistic inference model that departs from the standard practice in KD of using complex neural networks as teachers. The underlying idea is that an optimally parametrized sPDI model gives us access to a domain-specific probabilistic representation which (i) naturally produces geolocation highly dependent on BS locations, and (ii) can

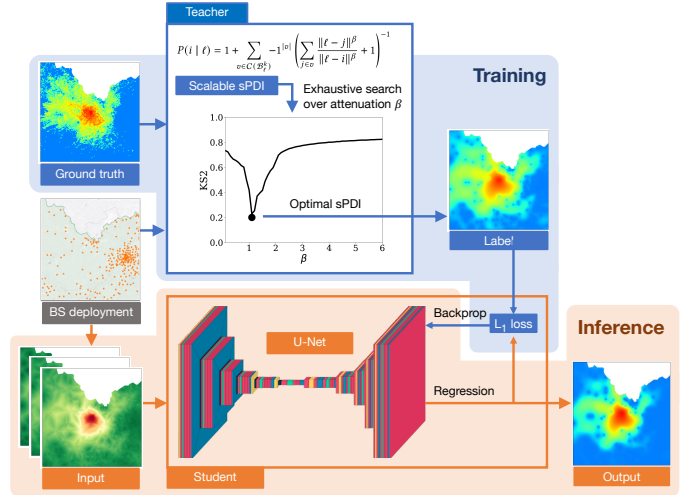


Fig. 3: Diagram of the DEEPMEND operation during training and inference. Once trained, the student can produce geolocation data from BS locations only (orange background).




Fig. 4: Metadata geolocation output by sPDI with different attenuation $\beta \in [0.5, 4]$ for the BS in Figure 1.

mimic the ground truth closely. The intuition is exemplified in Figure 4, for the same BS considered in Figure 1: the plots show the metadata geolocation probability $P(\ell|i)$ modeled by sPDI when its attenuation parameter β is varied from 0.5 to 4. We observe that the geolocation is too homogeneous over the whole geographical region for low β , and tends to the Voronoi cell for high β ; yet, some intermediate values, *e.g.*, $\beta = 2$ in the specific case of the considered BS, yield a geolocation that mimics well the ground truth in Figure 1a. We remark that the legacy recommended β value in this case for sPDI is 5 [31].

Optimal sPDI geolocation. Yet, the optimal β value of sPDI is BS-specific and depends on the usual combination of complex RF signal propagation and RAN configuration settings. Therefore, our teacher model has the objective of identifying, for each BS, the attenuation parameter producing an sPDI-generated geolocation that is the closest possible to the ground truth. By doing so, it outputs a geolocation that closely resembles the ground truth map for the target BS yet is generated solely from BS locations, *i.e.*, an ideal soft label.

As portrayed in Figure 3, our teacher achieves the goal above via an exhaustive search over the β space for each BS. More precisely, we run sPDI for all attenuation parameters between $\beta = 0.1$, which uniformly spreads the probability over a large region of thousands of km² around the BS, and $\beta = 6$, which converges to a traditional uniform geolocation over the Voronoi cell of the BS. For each BS i , we compare the




Fig. 5: Cost and accuracy of the scalable sPDI solution.

output of sPDI models parametrized with $\beta \in [0.1, 6]$ against the ground-truth geolocation using KS2 metric, and identify the β_i^* that minimizes the error for BS i . The teacher produces a soft label in the form of the geolocation generated by the sPDI model with the optimal attenuation β_i^* .

New scalable sPDI model. The exhaustive search for β_i^* for each BS i is very demanding from a computational viewpoint. While it is typical for teacher models to be complex, executing our proposed teacher with the original implementation of the sPDI model [32] for the thousands of BSs in our dataset would require months in a high-end computing cluster. We thus propose a new fast variant of sPDI as follows.

Given a tessellation of the target region into locations ℓ , and the set \mathcal{B} of BSs covering such region, the sPDI model defines the association probability of a UE in ℓ to the i -th BS as

$$P(i | \ell) = \int_0^\infty e^{-r} \prod_{j \in \mathcal{B} \setminus i} \left(1 - \exp\left(-r \frac{\|\ell - j\|^\beta}{\|\ell - i\|^\beta}\right) \right) dr, \quad (2)$$

where j are BSs other than i , and the attenuation parameter β indicates how quickly the signal weakens over space. Once (2) is solved, $P(\ell | i)$ can be derived using the Bayes theorem as per (1). A major disadvantage of the sPDI implementation in (1) is that the model entails a multiplication over all $j \in \mathcal{B} \setminus i$ for every location ℓ . In realistic scenarios where the number of BSs in \mathcal{B} is large (e.g., hundreds or more), the cost of computing the integral numerically makes our exhaustive-search teacher not viable even with high-end platforms.

In order to make sPDI scalable, we derive the analytical solution³ for $P(i | \ell)$ that avoids solving (2) numerically. Our novel closed-form solution of sPDI is expressed as

$$P(i | \ell) = 1 + \sum_{v \in \mathcal{C}(\mathcal{B}_\ell^k)} -1^{|v|} \left(\sum_{j \in v} \frac{\|\ell - j\|^\beta}{\|\ell - i\|^\beta} + 1 \right)^{-1}, \quad (3)$$

where $\mathcal{C}(\cdot)$ denotes the collection of all possible combinations of elements in the argument set and $-1^{|v|}$ is a sign factor that depends on whether the cardinality of v is odd or even. The equation is solved considering only the set \mathcal{B}_ℓ^k of k nearest neighboring BSs to each location ℓ which is efficiently computed in $\mathcal{O}(\log|\mathcal{B}|)$ via a KDTree data structure [38].

Cost-accuracy trade-off. Our proposed implementation of sPDI introduces a trade-off between efficiency and accuracy

³The derivation is omitted due to space limitations, but it relies on replacing the product inside the integral by the sum of all combinations over the set $\mathcal{B} \setminus i$, extracting the sum from the integral, and computing the sum of integrals.




Fig. 6: Example of the encoder-decoder input tensor channels (a)–(c) and their target soft label (d), for the BS in Figure 1.

when truncating (3) to a limited subset \mathcal{B}_ℓ^k of k neighboring BSs for each location. We explore this trade-off in Figure 5, which reports the results of (3) for different values of k .

Figure 5a presents the time required to compute $P(i | \ell)$ via (3) versus k : the fitted curve highlights an exponential growth of the computational complexity when the cardinality of \mathcal{B}_ℓ^k increases. Figure 5b depicts instead the relative Mean Square Error (MSE) of the geolocation $P(i | \ell)$ produced by (3), across values of k : as expected, larger k 's inherently produce better representations of the metadata geolocation. By juxtaposing the two plots in Figure 5, $k = 5$ emerges as a good operational point that grants both similar accuracy as models with higher k and bounded computational costs with respect to lower- k variants. As k reflects the number of BSs in proximity of a location that may have an impact on the association of the local mobile devices, it is interesting to observe that a value of 5 BSs is coherent with the typical number of BSs of a same operator covering a given location. Therefore, we employ $k = 5$ in our experiments in the rest of the document.

B. Student model

Our student is a deep neural network that performs a *pixel-to-pixel* (P2P) prediction of the geolocation probability $P(\ell | i)$ at each element ℓ of the regular grid used to represent the geographical area of the i -th BS. Importantly, we aim at a P2P prediction based only on input derived from BS locations; this starkly contrasts with works applying P2P to radio coverage, which need detailed terrain maps, building information, channel characteristics, and antenna configurations, modulation or code rates [24], [25], [39].

U-Net neural network. The core of the student model is a U-Net architecture [40]. We configure it with six encoding layers: the first layer has 32 convolutional filters, and each following one decreases the filters by a factor of two. The decoding part has a structure symmetric to that of the encoder. Each encoding or decoding layer uses a standard convolution with an kernel size $m = 5 \times 5$, followed by 2×2 max pooling or upsampling operation, respectively. All layers use ReLu activation functions, except for the final layer that uses a 1×1 convolution and linear activation to extract the output tensor. The choice of the student neural network is the result of extensive experiments with different architectures, e.g., U-Net, SDU-Net [41], and diverse hyper-parametrizations.

Input transformation. The U-Net is not fed with the raw BS locations, but with an original transformation of the BS positioning, which is more informative for the P2P

prediction. In particular, we use as input a 600×600 tensor of locations ℓ centered at the target BS and composed of three channels, which capture the distance between each pixel (*i.e.*, location) and its first, second, and third closest BS. The rationale for this choice is that these are the key information leveraged by the sPDI model in (2); hence, they carry strong P2P relationships with the target geolocation map that the models can learn. To further facilitate the encoder-decoder architectures to develop a connection between the input and the output tensor, a transformation is applied to the input channels $X'_{\ell,c} = 1 - \exp(-X_{\ell,c}/d_i(\ell))$. Here, $X_{\ell,c}$ is the initial value of pixel ℓ in the c -th channel, while $d_i(\ell)$ is the distance between pixel ℓ and the target BS i placed at the center of the map. This transformation renders the input of the encoder-decoder alike to the information used to calculate the metadata geolocation via (2), thus facilitating the training procedure. An example of the input tensor X is shown in Figure 6, along with the target soft label, for the BS considered in Figure 1.

Training. Abiding by KD principles, the loss is the L_1 distance between the output of the neural network and the soft label y provided by the teacher, represented as a 600×600 single-channel tensor with the values of the objective metadata geolocation. Since the target label often takes small values, we apply a logarithmic transformation $y' = -1/\log_{10}(y)$ to the output, and then re-normalize to a Probability Density Function (PDF); this lets the student model identify better locations with low geolocation probability and prevents the training process from getting stuck in local minima.

V. PERFORMANCE EVALUATION

We implement DEEPMEND using Tensorflow [42]. In all experiments, we train the model for 100 epochs, using the Adam optimizer with learning rate $1e^{-4}$ and batch size 8, on three NVIDIA A100 GPUs. In Section V-A, we evaluate the performance of DEEPMEND by conducting an n -fold cross-validation, with $n = 5$. In each instance, the full data set of BSs is split into training and test data with the 80:20 ratio commonly used in ML tasks [43]. That lets us study the overall performance of our approach in comparison with other solutions proposed in the literature. Then, in Section V-B, we explore how DEEPMEND generalizes to cases that are drastically different than those seen during the training phase, by training our model with a *leave-one-out* (LOO) strategy [43]. Specifically, we remove from the data set all the BSs in a specific (*i*) city or (*ii*) urbanization level, we train DEEPMEND with the remaining samples, and then probe the quality of its predictions for the city or urbanization level left out. This process shows how DEEPMEND can serve for new scenarios and unknown environments not seen during training.

A. Comparative analysis

Across our evaluations, we compare DEEPMEND against benchmarks that include the state-of-the-art solutions for mobile network metadata geolocation, as follows.

- **Voronoi** uniform geolocation over the plain Voronoi cell of a BS, which is the common approach widely adopted in the literature, *e.g.*, [1]–[14], [16], [44], [45].

Model	Geolocation quality (ks2)	Inference (ms)
Voronoi	0.704 ± 0.3047	< 1
VoronoiBoost	0.454 ± 0.2316	3
sPDI	0.501 ± 0.2956	$\sim 80,000$
Direct	0.422 ± 0.2737	21
DEEPMEND	0.304 ± 0.1944	21

TABLE I: Summary of comparative performance evaluation.




Fig. 7: Violin plot of the ks2 error between the ground-truth geolocation and the approximations of the models.

- **VoronoiBoost** is a data-driven model that scales and composes the original Voronoi cells to create a probabilistic geolocation for each individual BS [20].
- **sPDI** is the sPDI model [32] configured with the default attenuations in the work that introduced PDI [31].
- **Direct** is a deep learning approach that aims at directly learning to predict the ground-truth geolocation from ground truth data. Essentially, this model is the student portion of DEEPMEND, trained without the help of the teacher, *i.e.*, on the ground truth instead of the soft labels. This benchmark allows assessing the importance of adopting KD for metadata geolocation estimation.

The results of the cross-validation are summarized in Table I. For each model, the quality of the geolocation it generates is expressed as the median ks2 error (with deviation) with respect to the ground truth. Inspecting the ks2, DEEPMEND substantially outperforms all solutions adopted or proposed in the literature. Namely, the median error is cut by 56.80%, 33.03%, and 39.27% over Voronoi, VoronoiBoost, and sPDI.

A deeper look is provided in Figure 7, which reports the statistics of the ks2 error metric obtained with each model across all cross-validation iterations. Again, DEEPMEND emerges as a clear winner, and it is, in particular, the only model to dramatically cut low-quality geolocation: while all other approaches have 30%–60% of BSs with poor ks2 > 0.6 , our model reduces those BSs to 9%. Ultimately, these results demonstrate that DEEPMEND can drastically enhance the quality of metadata geolocation with respect to approaches available in the literature.

The ks2 attained by DEEPMEND in Table I is also 27.87% lower than that achieved by the Direct benchmark. This significant difference clearly proves that learning to predict geolocation directly from ground-truth data, while solely using BS locations as input, is a hard task. It also demonstrates how the KD strategy we propose is key to achieving more reliable metadata geolocation representations. To complement




Fig. 8: Distributions of KS2 between the ground-truth metadata geolocation of each BS and DEEPMEND (green) and Direct (purple) models, across urbanization levels.




Fig. 9: Visual comparison of metadata geolocation maps.

the results, we also provide: (i) PDFs of KS2 for DEEPMEND and Direct for each urbanization level, in Figure 8, which allow a comparison against those of the other approaches in Figure 2; and, (ii) visual samples of the geolocation maps generated by all models, juxtaposed to the ground truth, in Figure 9. These visualizations substantiate how DEEPMEND provides geolocation maps close to the ground truth, while competitors fail to do so.

B. Generalization

We now validate the capability of DEEPMEND to generalize to new scenarios in environments different than those used during training. To this end, we leverage the fact that the ground-truth data described in Section III-A covers large regions across France and thus encompasses tens of cities as well as urban, suburban, and rural areas. This lets us adopt a LOO strategy at the level of (i) cities and (ii) urbanization levels. Specifically, for LOO on cities, we train DEEPMEND on BSs located in all cities in our reference dataset except one

City	DEEPMEND	Voronoi-Boost	Direct	sPDI	Voronoi
Bordeaux	0.39 ± 0.171	0.54 ± 0.186	0.56 ± 0.253	0.64 ± 0.265	0.86 ± 0.230
Dijon	0.38 ± 0.114	0.61 ± 0.149	0.56 ± 0.171	0.68 ± 0.194	0.89 ± 0.143
Lille	0.30 ± 0.162	0.41 ± 0.192	0.41 ± 0.240	0.50 ± 0.267	0.64 ± 0.272
Lyon	0.47 ± 0.182	0.67 ± 0.221	0.74 ± 0.267	0.81 ± 0.278	0.92 ± 0.257
Mans	0.34 ± 0.144	0.63 ± 0.153	0.53 ± 0.179	0.66 ± 0.195	0.88 ± 0.181
Marseille	0.32 ± 0.123	0.51 ± 0.166	0.51 ± 0.218	0.51 ± 0.221	0.72 ± 0.215
Nantes	0.36 ± 0.162	0.54 ± 0.189	0.60 ± 0.245	0.73 ± 0.248	0.89 ± 0.219
Nice	0.45 ± 0.157	0.52 ± 0.194	0.55 ± 0.277	0.54 ± 0.274	0.77 ± 0.240
Orleans	0.38 ± 0.187	0.63 ± 0.210	0.54 ± 0.232	0.61 ± 0.220	0.85 ± 0.245
Paris	0.15 ± 0.173	0.23 ± 0.199	0.19 ± 0.245	0.21 ± 0.241	0.33 ± 0.275
Rennes	0.45 ± 0.212	0.61 ± 0.206	0.53 ± 0.237	0.70 ± 0.283	0.87 ± 0.275
Sainttienne	0.40 ± 0.150	0.60 ± 0.203	0.64 ± 0.239	0.51 ± 0.195	0.86 ± 0.230
Toulouse	0.47 ± 0.203	0.57 ± 0.220	0.60 ± 0.284	0.72 ± 0.294	0.89 ± 0.265
Tours	0.39 ± 0.167	0.62 ± 0.169	0.62 ± 0.218	0.68 ± 0.235	0.88 ± 0.210
Urban	0.31 ± 0.231	0.37 ± 0.260	0.33 ± 0.306	0.39 ± 0.323	0.53 ± 0.322
Suburban	0.34 ± 0.139	0.52 ± 0.163	0.48 ± 0.219	0.60 ± 0.244	0.83 ± 0.224
Rural	0.32 ± 0.113	0.52 ± 0.131	0.48 ± 0.163	0.61 ± 0.182	0.85 ± 0.153

TABLE II: Summary of LOO generalization experiments.

and then test the resulting model on the left-out city. Similarly, for LOO on urbanization levels, we train DEEPMEND on BSs in urban and suburban areas and then test the trained model on BSs in rural areas; we then repeat the experiment, leaving out also urban or suburban BSs. Overall, these settings ensure that DEEPMEND is tested in diverse environments from those seen during training, characterized by a different density and typology of buildings, infrastructures, and green areas.

The performance of DEEPMEND in these experiments is reported in Table II. In all tests, DEEPMEND produces much higher fidelity geolocation maps compared to the benchmarks, corroborating the resilience of our framework. Indeed, Voronoi and sPDI (which do not need training, hence are simply run for the BSs in the target LOO areas), VoronoiBoost or Direct all result in clear performance drops with respect to our model: DEEPMEND cuts KS2 errors by 10%–50% with respect to the second-best model in each case.

When looking at the performance of DEEPMEND, we observe that for urbanization level LOO tests, errors stay comparable to those seen in Section V-A, with a KS2 within 0.31–0.34. Instead, the geolocation quality becomes more heterogeneous in LOO tests across cities, with high variance depending on the left-out city and KS2 ranging from 0.15 to 0.47. We remark that DEEPMEND tests show discrepancies of training and test errors of 1%–4% that exclude the possibility of overfitting; we thus ascribe the result to the fact that cities are very diverse, and we may need a larger training dataset to capture correlations that are specific to one precise conurbation. Yet, even in LOO cities with the worst performance DEEPMEND yields superior performance to all other models.

VI. APPLICATION USE CASES

We demonstrate how DEEPMEND can benefit studies based on network metadata by showing that it improves the credibility of the evaluations of two applications: vRAN planning, and dynamic demographic density mapping. We highlight that, in both use cases, we use geolocation generated in a city *unseen during LOO model training*, as per Section V-B.

A. vRAN planning

We consider a vRAN environment that comprises multiple Radio Units (RUs) that execute basic operations, such as FFT, Cyclic Prefix, or P/S, and are densely deployed over a target region. The vRAN infrastructure then includes Distributed Units (DUs) that perform more demanding PHY/MAC functions,




Fig. 10: (a) Jain’s fairness index of the traffic load served by deployed ESs. (b) MAPE of the dynamic population estimate.

such as coding, modulation, or FEC. The DUs are found at far Edge sites (ES), where they dispose of substantial computational resources [46]. For a successful vRAN operation, the traffic demands at the RUs must be taken into account during the deployment of ESs, in order to optimally exploit the ES capacity, and reduce their energy footprint [47].

We use our metadata geolocation model to generate detailed geographical distributions of the average mobile demand and then exploit such maps for ES planning. Specifically, we first use DEEPMEND to predict geolocation maps for all BSs in a large metropolitan area not seen during training. We then perform a spatial mapping of the average traffic load recorded by the network operator at each BS in the target area, by: (i) distributing the traffic volume of each BS to locations ℓ according to the DEEPMEND geolocation probability; and, (ii) summing up at each ℓ the load contributed by all BSs. The outcome is a fine-grained map of the typical traffic demand generated within each location ℓ in the region, denoted by d_ℓ .

Now, let us formally define the ES planning problem. We assume a dense and regular deployment of RUs $r \in \mathcal{R}$, each serving one location ℓ with demand $d_r = d_\ell$. We aim at deploying a set \mathcal{S} of k total ESs, so as to balance the volume of traffic handled by each site $s \in \mathcal{S}$, while also minimizing the fronthaul latency (*i.e.*, geographical distance) to the RUs it serves. We represent the vRAN as a graph $G(\mathcal{R}, \mathcal{E})$, where an edge $(r, r') \in \mathcal{E}$ only exists if RUs r and r' serve adjacent locations. Then, the problem is formulated as

$$\min \sum_{(r,r') \in \mathcal{E}} x(r, r'), \quad \text{s.t.} \quad (4)$$

$$1 - \epsilon \leq \frac{\sum_{r \in \mathcal{R}} y(r, s) \cdot d_r}{\sum_{r \in \mathcal{R}} d_r / |\mathcal{S}|} \leq 1 + \epsilon, \quad \forall s \in \mathcal{S}, \quad (5)$$

$$\sum_{s \in \mathcal{S}} y(r, s) = 1, \quad \forall r \in \mathcal{R}, \quad (6)$$

$$x(r, r') \geq y(r', s) - y(r, s), \quad \forall (r, r') \in \mathcal{E}, \quad \forall s \in \mathcal{S}, \quad (7)$$

$$x(r, r') \geq y(r, s) - y(r', s), \quad \forall (r, r') \in \mathcal{E}, \quad \forall s \in \mathcal{S}, \quad (8)$$

where $x(r, r')$ and $y(r, s)$ are decision variables equal to one if edge (r, r') is cut by a partition, and if RU r is associated with ES s , respectively. The problem in (4) is solved via the KaFFPa heuristic [48], which partitions of the RUs across ESs, so that ES traffic loads are balanced within a margin ϵ , as per (5). ESs are placed at partition barycenters to reduce fronthaul latency.

Having deployed the k ESs based on DEEPMEND-predicted data, we assess the planning quality by computing

how balanced are the *actual* demands served by each ES. Thus, we compute the same traffic map as described before, but using the ground truth metadata geolocation of the network operator from Section III-A, and obtain the ground-truth traffic $d_r^* = d_\ell^*$ at RU r . Then, the actual load at ES s is $d_s^* = \sum_{r \in \mathcal{R}} d_r^* \cdot y(r, s)$, where we recall that $y(r, s)$ is derived using information supplied by DEEPMEND. We finally use Jain’s fairness index over all samples d_s^* , $s \in \mathcal{S}$ to evaluate the load symmetry across ESs, hence the planning quality. This experiment allows assessing how an ES deployment would perform in the field (which, in our case, coincides with the ground-truth traffic map), even if the deployment were driven by geolocation estimated via DEEPMEND.

Figure 10a shows the evolution of fairness versus the density of the ES deployment. A DEEPMEND-informed planning yields excellent fairness close to 1 for a small number ESs; the quality of the selected locations decreases for a larger k where the planning choices also grow, yet stays close to 0.95 even for dense deployments of hundreds of ESs. We then solve the same planning problem when the traffic d_ℓ is generated from the geolocation maps estimated Voronoi, VoronoiBoost, the legacy-parametrized sPDI model, and the Direct approach. Figure 10a reports how the quality of the ES deployment fairness deteriorates much quicker, implying that relying on existing geolocation strategies to evaluate ES planning can bias conclusions on the expected performance of the vRAN.

B. Dynamic demographic density

Estimating the population presence at fine spatial granularity over time is an open problem in demography, whose solution can improve, *e.g.*, transportation system design or urban policing. This is a challenging problem utterly different from surveying dwelling units as done in the population census: dynamic demographic density estimation aims to follow order-of-minute variations in the presence of people in urban areas, whereas census data only captures their static home locations. Mobile network metadata is emerging as a valuable source of information for models of real-time population density [1]. In this use case, we consider a multivariate regression model for the estimation of demographic density $p_\ell(t)$ at location ℓ and time t , which leverages the mobile traffic demand $d_\ell(t)$ measured at t in ℓ [49]. Formally,

$$p_\ell(t) = e^{k_1 \lambda(t) + k_2} d_\ell(t)^{k_3 \lambda(t) + k_4}, \quad (9)$$

where $\lambda(t)$ is the mean number of network events (*e.g.*, data sessions) per subscriber, while k_1 , k_2 , k_3 , and k_4 are constants. We set all parameters in (9) to empirical values provided in the original paper [49], and compute hourly traffic demand $d_\ell(t)$ with geolocation maps produced by DEEPMEND or the benchmarks, exactly as in Section VI-A.

We then use the model in (9) to generate hourly cartographies of the demographic presence in the target metropolitan area. We compare such presence to the ground-truth hourly population density, which we estimate again with (9), but using the realistic demands $d_\ell^*(t)$ generated from operator-provided geolocation data. A quantitative view of the model

quality is provided in Figure 10b, which shows the Mean Absolute Percentage Error (MAPE) obtained by comparing population density informed by the ground-truth versus model-based geolocation across one day. DEEPMEND yields an error around 15%. When repeating the same experiments with currently available geolocation methods, our approach yields errors 25% to 50% lower than the benchmarks. We conclude that our model can improve, e.g., urban policing based on metadata-driven dynamic population estimates.

VII. CONCLUSIONS

We propose DEEPMEND, a data-driven model that can predict high-resolution metadata geolocation maps from BS locations only. DEEPMEND hinges upon an original knowledge distillation design with suitable probabilistic inference and deep learning pixel-to-pixel models. It is accurate, computationally efficient, generalizes well to previously unseen environments, and largely improves existing geolocation methods. DEEPMEND has the potential to democratize access to reliable geolocation maps, improve the quality of research based on mobile network metadata, and contribute to open tools for producing official statistics from network metadata.

ACKNOWLEDGMENT

This work was supported by NetSense, grants no. 2019-T1/TIC-16037 and 2023-5A/TIC-28944, funded by Comunidad de Madrid; and by CoCo5G, grant no. ANR-22-CE25-0016, funded by the French National Research Agency. The work of S. Bakirtzis was supported by the Onassis Foundation and the Foundation for Education and European Culture.

REFERENCES

- [1] P. Deville *et al.*, “Dynamic population mapping using mobile phone data,” *PNAS*, 2014.
- [2] F. Batista *et al.*, “Uncovering temporal changes in europe’s population density patterns using a data fusion approach,” *Nat. Commun.*, 2020.
- [3] M. Lenormand *et al.*, “Comparing and modelling land use organization in cities,” *The Royal Society Open Science*, 2015.
- [4] M. Seppcher *et al.*, “Estimation of urban zonal speed dynamics from user-activity-dependent positioning data and regional paths,” *Transportation Research Part C: Emerging Technologies*, 2021.
- [5] M. De Nadai *et al.*, “The death and life of great italian cities: a mobile phone data perspective,” in *The WWW Conference*, 2016.
- [6] W. Chen *et al.*, “Impact of air pollution on human activities: Evidence from nine million mobile phone users,” *PLoS ONE*, 2021.
- [7] S. Mishra *et al.*, “Second-level digital divide: A longitudinal study of mobile traffic consumption imbalance in france,” in *The WWW Conference*, 2022.
- [8] J. E. Steele *et al.*, “Mapping poverty using mobile phone and satellite data,” *The Royal Society Interface*, 2017.
- [9] I. Ucar *et al.*, “News or social media? socio-economic divide of mobile service consumption,” *The Royal Society Interface*, 2021.
- [10] T. Y. *et al.*, “Mobile phone location data for disasters: A review from natural hazards and epidemics,” *Comput Environ Urban Syst*, 2022.
- [11] A. Vazquez Brust *et al.*, “Detecting areas of potential high prevalence of chagas in argentina,” in *The WWW Conference*, 2019.
- [12] S. Heroy *et al.*, “Covid-19 policy analysis: labour structure dictates lockdown mobility behaviour,” *The Royal Society Interface*, 2021.
- [13] U. Paul *et al.*, “Understanding traffic dynamics in cellular data networks,” in *IEEE INFOCOM*, 2011.
- [14] Z. Fang *et al.*, “Cellrep: Usage representativeness modeling and correction based on multiple city-scale cellular networks,” in *The WWW Conference*, 2020.
- [15] S. Mishra *et al.*, “Characterizing 5g adoption and its impact on network traffic and mobile service consumption,” in *IEEE Infocom*, 2024.
- [16] C. Zhang *et al.*, “Long-term mobile traffic forecasting using deep spatio-temporal neural networks,” in *MobiHoc*, 2018.
- [17] F. Ricciato *et al.*, “Towards a methodological framework for estimating present population density from mobile network operator data,” *Pervasive and Mobile Computing*, 2020.
- [18] Flowminder, “Flowkit.” <https://github.com/Flowminder/FlowKit>, 2023.
- [19] F. Ricciato and A. Coluccia, “On the estimation of spatial density from mobile network operator data,” *TMC*, 2021.
- [20] O. E. Martínez-Durive *et al.*, “Voronoiboost: Data-driven probabilistic spatial mapping of mobile network metadata,” in *SECON*, 2022.
- [21] V. D. Blondel *et al.*, “A survey of results on mobile phone datasets analysis,” *EPJ data science*, 2015.
- [22] J. Robinson *et al.*, “Assessment of urban-scale wireless networks with a small number of measurements,” in *ACM MobiCom*, 2008.
- [23] E. Alimpertis *et al.*, “City-wide signal strength maps: Prediction with random forests,” in *The WWW Conference*, 2019.
- [24] X. Zhang *et al.*, “Cellular network radio propagation modeling with deep convolutional neural networks,” in *KDD*, 2020.
- [25] S. Bakirtzis *et al.*, “EM DeepRay: An expedient, generalizable and realistic data-driven indoor propagation model,” *IEEE Trans. Antennas Propag.*, 2022.
- [26] H. Teerayut, *A study on urban mobility and dynamic population estimation by using aggregate mobile phone sources*. PhD thesis, Tokyo University, 2010.
- [27] F. de Meersman *et al.*, “Assessing the quality of mobile phone data as a source of statistics,” in *European Conference on Quality in Official Statistics*, Eurostat, 2016.
- [28] E. Graells-Garrido *et al.*, “Sensing urban patterns with antenna mappings: The case of santiago, chile,” *Sensors*, 2016.
- [29] O. Järv *et al.*, “Enhancing spatial accuracy of mobile phone data using multi-temporal dasymetric interpolation,” *International Journal of Geographical Information Science*, 2017.
- [30] J. Portela *et al.*, “Cellular network as a multiplicatively weighted voronoi diagram,” in *Consumer Communications and Networking*, 2006.
- [31] H. Zang *et al.*, “Bayesian inference for localization in cellular networks,” in *IEEE INFOCOM*, 2010.
- [32] V. A. Traag *et al.*, “Social event detection in massive mobile phone data using probabilistic location inference,” in *Conference on Privacy, Security, Risk and Trust*, 2011.
- [33] M. Tennekis *et al.*, “A bayesian approach to location estimation of mobile devices from mobile network operator data,” *Journal of Spatial Information Science*, 2021.
- [34] J. A. Peacock, “Two-dimensional goodness-of-fit testing in astronomy,” *The Royal Astronomical Society*, 1983.
- [35] G. Fasano and A. Franceschini, “A multidimensional version of the kolmogorov–smirnov test,” *The Royal Astronomical Society*, 1987.
- [36] C. D’Alessandro *et al.*, “La france et ses territoires,” Apr 2021.
- [37] G. Hinton *et al.*, “Distilling the knowledge in a neural network,” in *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [38] S. Maneewongvatana and D. M. Mount, “On the efficiency of nearest neighbor searching with data clustered in lower dimensions,” in *International Conference on Computational Science*, 2001.
- [39] A. Gupta *et al.*, “Machine learning-based urban canyon path loss prediction using 28 ghz manhattan measurements,” *IEEE Transactions on Antennas and Propagation*, 2022.
- [40] O. Ronneberger *et al.*, “U-net: Convolutional networks for biomedical image segmentation,” in *Med Image Comput Comput Assist Interv*, 2015.
- [41] S. Wang *et al.*, “U-net using stacked dilated convolutions for medical image segmentation,” *arXiv preprint arXiv:2004.03466*, 2020.
- [42] M. Abadi *et al.*, “Tensorflow: a system for large-scale machine learning,” in *USENIX OSDI*, 2016.
- [43] I. Goodfellow *et al.*, *Deep learning*. MIT press, 2016.
- [44] N. Pokhriyal and D. C. Jacques, “Combining disparate data sources for improved poverty prediction and mapping,” *PNAS*, 2017.
- [45] N. Oliver *et al.*, “Mobile phone data for informing public health actions across the covid-19 pandemic life cycle,” *Science Advances*, 2020.
- [46] 3GPP, “Release description; release 15,” Technical report (TR) 21.915, 3rd Generation Partnership Project (3GPP), 06 2019. Version 1.1.0.
- [47] R. Singh *et al.*, “Energy-efficient orchestration of metro-scale 5g radio access networks,” in *IEEE INFOCOM*, 2021.
- [48] C. S. Pan *et al.*, “Very large graph partitioning by means of parallel dbms,” in *Advances in Databases and Information Systems*, 2013.
- [49] G. K. *et al.*, “Estimation of static and dynamic urban populations with mobile network metadata,” *TMC*, 2018.