



This research is part of the 'LINC project'. This project is co-financed by the European Regional and Development Fund through the Urban Innovative Actions Initiative. The LINC project consists of a larger consortium led by Gate21. The consortium includes Municipality of Albertslund, Municipality of Gladsaxe, Nobina Danmark A/S, IBM Danmark ApS, Roskilde University (RUC) and The Technical University of Denmark (DTU). The project is funded by project partners and the EU programme Urban Innovative Actions (UIA), which is supporting the project with 25 million DKK.

LARGE SCALE PASSENGER DETECTION WITH SMARTPHONE/BUS IMPLICIT INTERACTION AND MULTISENSORY UNSUPERVISED CAUSE-EFFECT LEARNING

A PREPRINT

Valentino Servizi Department of Technology, Management and Economics Technical University of Denmark (DTU) valse@dtu.dk

Dan R. Persson Department of Applied Mathematics and Computer Science DTU

Per Bækgaard

Department of Applied Mathematics

and Computer Science

DTU

Francisco C. Pereira Department of Technology, Management and Economics DTU

Jeppe Rich

Department of Technology,

Management and Economics

DTU

Hannah Villadsen Department of People and Technology Roskilde University Denmark

> Otto A. Nielsen Department of Technology, Management and Economics DTU

February 20, 2022

ABSTRACT

Intelligent Transportation Systems (ITS) underpin the concept of Mobility as a Service (MaaS), which requires universal and seamless users' access across multiple public and private transportation systems while allowing operators' proportional revenue sharing. Current user sensing technologies such as Walk-in/Walk-out (WIWO) and Check-in/Check-out (CICO) have limited scalability for large-scale deployments. These limitations prevent ITS from supporting analysis, optimization, calculation of revenue sharing, and control of MaaS comfort, safety, and efficiency. We focus on the concept of implicit Be-in/Be-out (BIBO) smartphone-sensing and classification.

To close the gap and enhance smartphones towards MaaS, we developed a proprietary smartphonesensing platform collecting contemporary Bluetooth Low Energy (BLE) signals from BLE devices installed on buses and Global Positioning System (GPS) locations of both buses and smartphones. To enable the training of a model based on GPS features against the BLE pseudo-label, we propose the Cause-Effect Multitask Wasserstein Autoencoder (CEMWA). CEMWA combines and extends several frameworks around Wasserstein autoencoders and neural networks. As a dimensionality reduction tool, CEMWA obtains an auto-validated representation of a latent space describing users' smartphones within the transport system. This representation allows BIBO clustering via DBSCAN. We perform an ablation study of CEMWA's alternative architectures and benchmark against the best available supervised methods. We analyze performance's sensitivity to label quality. Under the naïve assumption of accurate ground truth, XGBoost outperforms CEMWA. Although XGBoost and Random Forest prove to be tolerant to label noise, CEMWA is agnostic to label noise by design and provides the best performance with an 88% F1 score.

arXiv:submit/4174401 [cs.LG] 20 Feb 2022

1 Introduction

Tracking passenger movements through the public transport network, seamlessly and without direct human interaction, requires accurate models and methods to discriminate between passengers that are using the public transport network and anyone else outside the transport network. While the accurate solution of such an implicit Be-In/Be-Out (BIBO) classification problem [Narzt et al., 2015], is directly relevant as a mean to collect important data from the public transport system, e.g. Check-in/Check-out or Walk-in/Walk-out statistics, it is relevant for other areas as well. This includes as an example, the tracking of persons entering buildings to comply with safety measures and the registration, and tracking of people in supermarkets to support crew management in different parts of the supermarket. However, tracking of public transport users represent a more complex problem in that buses and passengers move in space-and time. As a result, we will argue that the ability to provide robust solutions for public transport applications is a stepping-stone for these other relevant applications.

Solving the before mentioned classification problem is important for several reasons. Firstly, on the very practical side it provides a means to collect valuable data about passenger flows that would otherwise have been lost for users paying by cash, or accidentally traveling without checking-in. Secondly, it would enable context-aware surveying and services while lifting the burden of explicit interaction from passengers. Thirdly, for planning optimal departure times and routes of a trip through the public network, it would support personalized dynamic recommendations.

In a wider perspective the presented methodology can be seen as an important component in Mobility-as-a-service (MaaS) systems. MaaS combines multiple transport modes as transport services–e.g., car, bus, bike, scooter–offered through a single interface, and paid with the same unique subscription, as the media contents on "Netflix" [Hietanen, 2014, Hensher and Mulley, 2021]. Hence, MaaS is essentially *"a data-driven, user-centered paradigm, powered by the growth of smartphones"* [Goodall et al., 2017]. Regardless from the perspective, MaaS ultimate goal is to enable a door-to-door public service, attractive for the passengers, and competitive with, e.g., privately owned cars. In this context, the ability to accurately track passengers while traveling would underpin the efficient capacity planning for a dynamic, responsive, and intelligent public transport paradigm.

In the MaaS context, smartphone-based automatic fare collection systems (AFCS) with BIBO could allow the integration of public service ticketing, automatic price calculation, and a fair cost split across multiple operators. The latter point includes emerging providers of, e.g., car- and bike-sharing services. Compared to CICO and WIWO, BIBO offers at least two advantages: (i) public transit increased comfort for passengers [Wirtz and Klähr, 2019]; and (ii) operational integration mostly software, with a negligible impact on new physical infrastructure. The second point means potentially lower access barriers for emerging transport service providers to MaaS. For the first, we refer to the passengers increased comfort with the term ticketless. Ticketless identifies the perspective of a system ability to flexibly adapting the transport service bill to the user's journey(s) across multiple service providers, as opposed to the perspective of multiple tickets necessary from multiple service providers, for the same journey.

From the Big Data perspective, handling this binary classification problem with supervised machine learning methods presents the following challenges:

- 1. Controlling noise in the labels:
- 2. Operating a sustainable labels collection cost;
- 3. Minimizing the impact of sensors and data collection on the battery; and
- 4. Minimizing the users' privacy exposure.

These challenges involve the service operator's perspective in the first case and the smartphone user's perspective in the others.

Although from a ticketing perspective there should be no noise, thus one should only be charged when he or she uses a transport service, when using tickets as labels to train machine learning algorithms, the assumption of possible undetected ticketing errors from both sides–passenger and service provider–seems more than reasonable.

Mining transport behavior from smartphones data relies, among other sensors, on Global Positioning System (GPS), Inertial Navigation System (INS), and Bluetooth Low Energy (BLE) signal[Servizi et al., 2021a]. In urban areas, where 80% of public transport demand occurs [Baescu and Christiansen, 2020] (e.g., in Denmark), the classification of sensors' observations is complex. With GPS, any transportation mode looks the same due to a combination of factors, such as GPS errors in urban canyons, proximity between pedestrians and buses, and vehicles' low speeds in congested traffic [Cui and Ge, 2003]. With INS, multiple habits, each corresponding to whether one carries a smartphone, e.g., in the pocket or the bag, determine different sensors patterns [Wang et al., 2019]; the integral of any noise included in the sensors' signal, in addition, leads to often unmanageable error drifts [Foxlin, 1996]. The BLE signal, which is

extensively studied for indoor tracking, presents an excellent potential for proximity sensing and battery efficiency [Bjerre-Nielsen et al., 2020]. However, smartphones' signal records of BLE devices in proximity suffer from signal gaps [Malmberg, 2014]; a higher spatial density of BLE devices allows good indoor-tracking performance, but such a density is not scalable at a city scale. In contrast, GPS and INS scaling potential correspond to a heavy impact on the smartphones' battery [Servizi et al., 2021a]. In the first case, the sensor is directly responsible for the energetic consumption. In the second case, the sensors' energy consumption is sustainable as long as the signals are classified online within the smartphone. Yet, due to the high sampling rate necessary for achieving acceptable classification performance, > 20 Hz, data consumption outside the smartphone would imply high network energy consumption for data transfer [Servizi et al., 2021a]. In the assumption of training a supervised machine learning algorithm with high-quality labels, BIBO binary classification in the urban context seems a difficult task. When labels' quality degrades, we face another limitation as classifiers' performance can be highly biased-consequently, decisions would be based on scores looking high when they are low in reality and vice-versa [Servizi et al., 2021b]. To overcome the limitations mentioned above, in this work, we rely on a unique dataset collected during three months of autonomous buses' operations across a local public network in Denmark. The dataset includes the GPS and BLE trajectories collected from buses and passengers' smartphones through a proprietary smartphone-sensing platform, including 300 BLE devices installed in buildings near the bus network, in the buses, and at bus stops. Another set of the data provides high-quality ground truth collected by users that followed precise instructions on individual sequences of origins and destinations within the bus network, along specific routes [Shankari et al., 2020].

1.1 Literature Review

The solution we propose for the BIBO classification problem involves the implicit interaction of passenger smartphones, buses, and bus-network [Servizi et al., 2021b, Narzt et al., 2015]. Therefore, it falls within the intersection of several disciplines converging around smartphone-based travel surveys and smartphones indoor tracking with BLE network interaction. In the first case, leveraging smartphone onboard sensors, we are interested in the limitations of the methods for mode detection in general and bus detection in particular [Wirtz and Klähr, 2019]; in the second case, we are interested in how to deal with BLE signals [Servizi et al., 2021b].

The literature on mode detection from smartphones data is pervasive. GPS and INS sensors are the most used also to provide location- and person-agnostic mode classification. GPS and INS systems generate very different trajectories. The first system provides a geospatial time series with a sampling rate ≥ 1 Hz [Servizi et al., 2020, Dabiri and Heaslip, 2018]; the second system, a three-dimension time series along the three axes of the smartphone's reference frame, and a sampling rate ≥ 20 Hz [Servizi et al., 2021a, Cornacchia et al., 2017]. To prepare the data for the classification, the steps one follows to clean and segment these trajectories differ too. However, the best-performing classification methods consist of two main groups. The first group includes supervised methods, such as decision trees, random forest, and XGBoost [Koushik et al., 2020]; the second group has various configurations of artificial neural networks (ANN), both supervised and semi-supervised. Unsupervised methods based on clustering are applied directly to features extracted from GPS and INS, but their performance seems below the supervised and semi-supervised methods mentioned above. The blooming literature on both GPS- and INS-based mode detection proposes very effective methodologies, equally accurate when datasets include urban and outskirt areas and multiple transportation targets [Servizi et al., 2021a]. However, at low speeds, state-of-the-art INS-based online classifiers available on the leading smartphone operation systems seem unable to discriminate between bus and walk mode. In contrast, GPS and BLE classifiers show higher performance [Servizi et al., 2021b].

Among the studies focusing on mode detection and public transportation, specifically buses, the most promising are considering the interaction between users and the transport network. This interaction could be expressed as the time series of the distances between each point of a smartphone's GPS trajectory and each point of interest (PoI) extracted from the infrastructure mapped on GIS [Semanjski et al., 2017]. The classification could be point-based, thus relying on short segments. Another approach, which we define segment-based [Servizi et al., 2021a], could look at longer trip segments and the periodicity of stops typical of any bus operation [Zhang et al., 2011]. However, while the first approach suffers the limitation from the GPS error in dense urban areas, the second approach seems ineffective for short trips.

Literature focusing on BLE and WiFi signals-both based on the same communication frequency and protocols sharing some similarities-converges between indoor tracking and mode detection. The traditional methodologies leverage the Friis equation, and the trilateration [Kotanen et al., 2003, Subhan et al., 2011]. However, machine learning methods such as random forests and Gaussian processes are effective in BLE or WiFi fingerprint classification, and spatial signal mapping [Chen et al., 2015, Subhan et al., 2013, Pérez Iglesias et al., 2012]. To allow optimal BIBO sensing and classification with BLE devices, we find no clear contributions on the minimum spatial density of BLE devices, nor how to cover the scale of a city [Servizi et al., 2021b]. Therefore, we rely on literature about indoor tracking [Yassin

et al., 2017] and preliminary BIBO experiments with BLE signals [Servizi et al., 2021b], suggesting that BLE devices installed in buses and bus stops could offer a coverage sufficient for classification. Consequently, such a configuration would have the potential to cover the entire city at a reasonable cost.

The parallel growth of computation power and data volume kept in check the tradeoff between computational capacity and classification performance. On the one hand, Computation Processing Units (CPU) and Graphical Processing Units (GPU) have created sizeable extra computation potential. On the other hand, the pursuit of better accuracy leveraging, for example, the pervasive introduction of cheap sensors and rich Geographic Information Systems (GIS), immediately absorbed this additional capacity. Overall, transportation mode classifiers deployed on data from urban and densely populated areas did not increase their performance proportionally with the data consumption. Therefore, statistical methods developed before the Big Data paradigm [Schuessler and Axhausen, 2009], and machine learning methods developed after [Koushik et al., 2020], may still compete. A factor emerging from the literature is that methods still depend heavily on labels. Even though some semi-supervised configuration of artificial neural networks exists in this field and reduces the need for labels in the classifier's training phase, filtering a subset of high-quality labels from Big dataset is still very challenging and hardly scalable. For example, continuous disruptions of transport operations due to roadwork or special events would also disrupt any classifier trained with labels that no longer reflect the transport network [Petersen et al., 2021]. Even in the assumption of operations stability, the impact of flipping and overlaying labels-potentially present due to human collection errors-seems still critical. Supervised classifiers deployed on time series, e.g., for the BIBO task, could deliver biased classifications and threaten the system's sustainability at scale. The problem deserves more attention in this field, and for time series requires at least the same attention granted to independent and identically distributed data. Systematic studies and appropriate methodologies in the second case exist, such as for image classification. However, for time series classification these contributions are only partially applicable. Furthermore, existing preliminary studies about the impact of flipping labels on time series classification show that severe bias on the measurements of these classifiers' performance is present when just 10% of the labels are wrong. In such a case, although the classifiers might be resilient to labels' noise, analysts and practitioners would base their decisions on a biased performance evaluation, simply because the error rate in human validated labels is unknown [Servizi et al., 2021b].

1.2 Contribution of the Paper

This paper focuses on the combined use of GPS and BLE signals for unsupervised autovalidated BIBO classification of bus passengers. Representing the user via the smartphone and the bus via a BLE device, we use sensors signals as pseudo labels to learn discriminating when a user is inside (BI) or outside (BO) the bus.

The central intuition is that when the user is inside the bus (BI) the distance between smartphone and bus should be close to zero, and the proximity to BLE devices installed in the bus would cause the highest signal strength. Vice-versa, when the user is outside the bus (BO), the considerable distance between the user and the BLE device should cause the lowest signal strength or no signal at all.

To learn the cause-effect relationship between smartphone-bus proximity and BLE signal strength, we implement two parallel Wasserstain Autoencoders (WAE). One learns how to reconstruct the time series of the BLE signal (effect) given the smartphone-bus proximity (cause). Given the BLE signal strength (effect), the other learns to rebuild the smartphone-bus distance (cause). We define this configuration as a cause-effect multi-task Wasserstein Auto-encoder (CEMWA). From the unsupervised training of this CEMWA, we learn to reduce the description of the interaction between passengers and buses to only four dimensions. In this 4-dimensional latent space, the observations self-organize such that discrimination between BI and BO classes is possible through unsupervised clustering with Density-based spatial clustering of applications with noise (DBSCAN).

CEMWA combines and extends the following frameworks. (i) Split-brain Auto-encoder configuration by Zhang et al. [2016]; (ii) Deep clustering for unsupervised learning [Caron et al., 2018]; (iii) Multi-task formulation of the objective function by Kendall et al. [2018]; (iv) Maximum Mean Discrepancy (MMD) formulation of the objective function for generative models by Gretton et al. [2008]; and (v) MMD extension to Wasserstein Auto-encoders by Tolstikhin et al. [2017].

The resulting architecture solves the scalability problem related to noise in labels. We perform an ablation study including traditional WAE architectures and supervised methods. Results show that our unsupervised classifier solves the negative impact of the label-induced bias affecting supervised classifiers. Moreover, the architecture we propose embodies a solution for signal data imputation, which is generally a critical and separate step necessary to perform good classification. Finally, since the method relies only on the interaction between smartphone and bus, temporary or permanent disruptions of the network would not affect the classification task.



Figure 1: Cause-effect Multi-task Wasserstein Auto-encoder (CEMWA) independent cross-reconstruction of X_1, X_2 minimizing (7) and clustering of the resulting latent space, 5028 parameters.



Figure 2: Multi-task Wasserstein Auto-encoder (MWA) independent reconstruction of (X_1, X_2) minimizing (3), with $c = \mathcal{L}_{WAE}$ and clustering of the resulting latent space, 5028 parameters.



Figure 3: Wasserstein Auto-encoder (WA) reconstruction of $X = (X_1, X_2)$ minimizing (1) and clustering of the resulting latent space, 4932 parameters.

2 Methods and Materials

This section presents a number of frameworks supporting our goal of substituting ordinary labels for training supervised or semi-supervised artificial neural networks specialized in processing GPS signal. Three are the main steps behind the intuition. Firstly, instead of labels we leverage an independent sensor time-series–BLE–for representation learning of cause-effect relationship between GPS and BLE. Secondly, to avoid confounding correlations between the two sensors' signals, we design and fine-tune a specific encoder-decoder architecture based on a general formulation of regularized auto-encoders. Lastly, with DBSCAN, we turn into classes the representations learned via independent sensors time-series–GPS and BLE.

Following the notation of Tolstikhin et al. [2017], we identify sets with calligraphic letters (i.e. \mathcal{X}), random variables with capital letters (i.e. X), and values with lower case letters (i.e. x).

Let $X \in \mathbb{R}^{t \times d}$ be the tensor describing the smartphone/bus interaction, in a time window of t observations, which d independent feature channels express such that: $X_1 \in \mathbb{R}^{t \times d_1}$ represents the channels deriving from the GPS sensors; $X_2 \in \mathbb{R}^{t \times d_2}$, from the BLE devices network; where $(X_1, X_2) = X$ and $\mathcal{D}_1 \bigcup \mathcal{D}_2 \subseteq \mathcal{D}$, with $|\mathcal{D}| = d$.

We would like to learn a representation for X solving the prediction problem $\hat{X} = (\hat{X}_1, \hat{X}_2)$, where $\hat{X}_1 = \mathcal{F}_1(X_2)$, and $\hat{X}_2 = \mathcal{F}_2(X_1)$. \mathcal{F}_1 learns the cause-effect relationship between smartphone-bus proximity and BLE signal strength, while \mathcal{F}_2 learns the inverse cause-effect relationship of the same interaction between smartphone and bus.

 \mathcal{F} represents a class of non-random generative Encoder/Decoder models determinalistically mapping input points to the latent space with a convolutional neural network (CNN) via Encoder, and latent codes to output points with a transpose CNN via Decoder. To learn \mathcal{F} , we minimize the Wasserstein optimal transport cost (1) between the true-unknown data distribution P_X and the latent variable model P_G specified by the prior distribution P_Z of latent codes $Z \in \mathcal{Z}$ and the generative model $P_G(X|Z)$ of the data points $X \in \mathcal{X}$ given Z [Tolstikhin et al., 2017]. (1) shows that while the decoder pursues the encoded training examples reconstruction at the minimal cost c, the encoder pursues two conflicting goals at the same time: (i) Match the encoded distribution Q_Z to the prior distribution P_Z , where $Q_Z := \mathbb{E}_{P_X} [Q(Z|X)]$ (ii) Ensure that the latent representation for the decoder allows accurate reconstruction of the encoded training examples.

In this two steps procedure, first Z is sampled from a fixed distribution P_Z on a latent space \mathcal{Z} , and then Z is mapped to $\hat{X} = G(Z)$ for a given map $G : \mathcal{Z} \to \mathcal{X}$, where $\hat{X} \in \mathcal{X} = \mathbb{R}^{t \times d}$.

$$\mathcal{L}_{WAE}(P_X, P_G) := \inf_{Q(Z|X) \in \mathscr{Q}} \mathbb{E}_{P_X} \mathbb{E}_{Q(Z|X)} \left[c(X, G(Z)) \right] \\ + \lambda \cdot \mathcal{D}_Z(Q_Z, P_Z),$$

$$\lambda > 0$$
(1)

This task formulation extends the Split-brain Autoencoder proposed by Zhang et al. [2016]. We share the intuition, and the goal of achieving a representation containing high-level abstraction and semantics of the smartphone-bus interaction registered independently by GPS and BLE sensors. In contrast with Zahng, we aim at learning the cause-effect function and its inverse, separately, and not just merely as a "pretext". However, to keep up with the Big Data scale, Zhang approach brings some limitations with the objective function in Eq. (2): (i) For weighting the multi-task cost \mathcal{O} , Zhang introduces the hyperparameter $\hat{\lambda}$ that requires a dedicated optimization process. (ii) To learn cause-effect relationship and its inverse, we do not want include the full signal $c((\mathcal{F}_1(X_2), \mathcal{F}_2(X_1)), X)$ in the multi-task objective function \mathcal{O} . (iii) The use of a classical unregularized auto-encoder, which minimizes only the reconstruction cost c, between X and \hat{X} , prevents from yielding full advantage of representation learning for this problem, facilitating model over-fitting instead of generalization power.

$$\mathcal{O} = \arg \min_{\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{F}} [\hat{\lambda} \cdot c(\mathcal{F}_2(X_1), X_2) \\ + \hat{\lambda} \cdot c(\mathcal{F}_1(X_2), X_1) \\ + (1 - 2 \cdot \hat{\lambda}) \cdot c((\mathcal{F}_1(X_2), \mathcal{F}_2(X_1)), X)],$$

$$\hat{\lambda} \in [0, \frac{1}{2}]$$

$$(2)$$

In the following sections we can now look at how we extended Zhang's work to cover both of the aforementioned limitations and enable clustering.

2.1 Extension Towards Multi-task Self-learned Cost Weights

In a multi-task setting, Kendall shows that when tasks uncertainty depends on its unit of measure, homoscedastic uncertainty is an effective bias for weighting multiple losses [Kendall et al., 2018]. This fits exactly with our problem, where the proximity between smartphone and bus is measured in meters on one hand, and in Received Signal Strength Indicator (RSSI) on the other hand. With $\hat{X}_1 = \mathcal{F}_1(X_2)$ and $\hat{X}_2 = \mathcal{F}_2(X_1)$, where $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{F}$, (3) represents the multi-task loss formulation for our problem, according to Kendall. The main difference between (2) and (3) is that in the second case the two parameters can be "learned" leveraging the ANN back propagation algorithm while learning \mathcal{F} parameters, during the training phase. When training on large datasets, this is an advantage.

$$\mathcal{O} = \arg \min_{c} \left[\frac{1}{2\sigma_{1}^{2}} \cdot c(\hat{X}_{1}, X_{1}) + \frac{1}{2\sigma_{2}^{2}} \cdot c(\hat{X}_{2}, X_{2}) + \ln \sigma_{1} + \ln \sigma_{2} \right]$$
(3)

2.2 Extension towards regularized auto-encoder

WAE represent a class of generative models resting on the optimal transport cost derived from Villani [2003] and expressed in (1). This class underpins our extension: In contrast to Zhang work [Zhang et al., 2016], which studies the unregularized cost c, such as regression and cross-entropy, we include to the regression cost a regularization term, i.e., the maximum mean discrepancy (MMD) $D_Z = MMD_k(P_Z, Q_z)$. (4) expresses the MMD, where $k : Z \times Z \to \mathbb{R}$ is a positive-definite reproducing kernel, and \mathcal{H}_k is the reproducing kernel Hilbert space (RKHS) of real-valued functions mapping Z to \mathbb{R} [Gretton et al., 2008].

Similarly to variational auto-encoders (VAE) [Kingma and Welling, 2013], this WAE-MMD formulation uses artificial neural networks (ANN) to parametrize encoder and decoder. However, to allow back-propagation throughout decoder

and encoder, the re-parametrization trick [Kingma and Welling, 2013] "forces Q(Z|X = x) to match P_Z for all the different samples x drawn from P_X . In contrast, WAE forces the continuous mixture $Q_Z := \int Q(Z|X)dP_X$ to match P_Z " [Tolstikhin et al., 2017]. Consequently, WAE allow a better organization of the latent space which we leverage for clustering. Compared to alternative formulations of the penalty term, such as the Generative Adversarial Networks [Makhzani et al., 2015] (GAN), or in general the WAE-GAN [Tolstikhin et al., 2017], where \mathcal{D}_Z in (1) is the Jensen-Shannon Divergence, the literature shows slightly better reconstruction performance for \hat{X} but at the heavy cost of an additional network and possibly complex and multi-modal distributions for P_Z . Since our problem is simple in principle, we opt for simplicity, thus for MMD.

$$MMD_{k}(P_{Z}, Q_{z}) = || \int_{\mathcal{Z}} k(z, \cdot) dP_{Z}(z) - \int_{\mathcal{Z}} k(z, \cdot) dQ_{Z}(z) ||_{\mathcal{H}_{k}},$$

$$(4)$$

If k is characteristic¹ MMD represents a divergence measure [Sriperumbudur et al., 2011].

We try both the alternative kernels k proposed for Wasserstein auto-encoders (WAE) [Tolstikhin et al., 2017]: Radial basis function kernel (RBF) (5); and Inverse multiquadratics kernel (6).

$$k^{RBF}(z,\tilde{z}) = e^{\frac{-||\tilde{z}-z||_2^2}{\hat{\sigma}_k^2}}$$
(5)

$$k^{IMK}(z,\tilde{z}) = \frac{C}{C + ||z - \tilde{z}||_2^2}$$
(6)

The resulting architecture consists of two independent encoder/decoder maps $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{F}$ such that $\hat{X}_1 = \mathcal{F}_1(X_2)$ and $\hat{X}_2 = \mathcal{F}_2(X_1)$. Each map's encoder consists of 1D-Convolutions; 1D-Transpose-Convolutions for the decoder. As described in Fig. 1, maps are learned using back-propagation to minimizing the multitask formulation of our objective function (7), where we set $c = ||X - \hat{X}||_2^2$ and $D_Z = MMD_k$. To find optimal relative weights between tasks, we leverage the same back-propagation algorithm.

$$\mathcal{O}_{WAE} = \arg \min_{\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{F}} \frac{1}{2\sigma_1^2} \cdot \mathcal{L}_{WAE}(\mathcal{F}_2(X_1), X_2) + \frac{1}{2\sigma_2^2} \cdot \mathcal{L}_{WAE}(\mathcal{F}_1(X_2), X_1) + \ln \sigma_1 + \ln \sigma_2$$
(7)

2.3 Extension of Deep Clustering Architecture

To allow unsupervised classification of images, Caron et al. proposes a straight ANN predicting cluster assignment as pseudo-labels [Caron et al., 2018], and iterate between clustering with k-means [Likas et al., 2003] and back-propagation to update the network's weights after the cluster assignment. The intuition is that clustering provides and alternative and meaningful reference to labels. Therefore, the loss function is computed against clusters instead of known labels. However, since we collect two independent measure of the same event, by design, we tweak the process using these two signal as reciprocal pseudo-labels instead. When back-propagation converges, we perform clustering of data representation on the latent space with DBSCAN [Khan et al., 2014]. Fig. 1, 2 and 3 show the architectures tested within our ablation study: the first leverages the known cause-effect relationship between GPS and BLE signal; the second, the multi-task independent reconstruction of the two signals; the last shares parameters within the same network, to reconstruct a tensor where multiple channels contain each available signal.

2.4 Final Model Formulation

Fig. 1 presents the final structure of our CEMWA model, resulting from the Split-brain's architecture extensions described in Sec. 2.1, 2.2 and 2.3.

¹Given $k : \mathcal{Z}^+ \to \mathbb{R}$, k is injective, \mathcal{Z}^+ is positive and represents the set of probability measures on \mathcal{Z}^+

We will argue as follows: (i) CEMWA has the ability of learning the cause-effect relationship between GPS and BLE signals recording smartphone-bus interactions. (ii) Learning such a relationship allows the exposure of self-validated features characterizing the BIBO status of users with respect to buses. (iii) These self-validated features allow unsupervised classification of users trajectories, where smartphones identify users and BLE devices identify buses. (iv) Alternative unsupervised architectures leveraging the correlation instead of cause/effect between the GPS and BLE signals—such as those described in Fig. 2 and 3— are unable to to perform self-validated unsupervised BIBO classification. (v) In case of labels noise, CEMWA significantly outperforms the most accurate supervised classifiers, such as random forest or XG-boost (extreme gradient boosting). (vi) Regardless of the classification performance, CEMWA embodies both a data imputation and a validation mechanism, while supervised classifiers or alternative unsupervised architectures should rely on dedicated processes, such as an exponential weighted moving average for BLE or GPS imputation [Osman et al., 2018], and user validation for BIBO labels [Servizi et al., 2021b,a].

To substantiate our hypotheses through the following experiments, consistently, we designed and deployed a specific sensing architecture, and collected high quality ground truth.

2.4.1 Ground truth collection, data cleansing, and preparation

CEMWA's architecture mirrors the smartphone sensing platform we designed and deployed to track the activity of three autonomous buses operating an experimental public service in Denmark, between two extremes of the Lyngby campus where the Technical University of Denmark is located.

During operations these buses are tracked via GPS available from the bus telemetry, while test passengers recruited for the experiment are tracked via smartphones. The sensing platform collected GPS signals that both smartphones and buses generate. GPS collection was strictly limited around the operations area using a geo-fence [Almomani et al., 2011]. In the same area, we deployed 300 BLE devices: one on each bus and bus stop, plus one at the entrance/s of each building in the campus.

To become a test passenger, each user provided explicit agreement to terms and conditions presented in compliance with the General Data Protection Regulation². The sensing platform supports both Android and iOS devices, and the Apps are published on GooglePlay³ and App Store⁴ respectively. This project is a social science study, includes data and numbers only, is not a health science project, and does not include human biological material nor medical devices. Consequently, in Denmark, where the data collection took place, the Health Research Ethics Act provides a dispensation for notification to any research ethics committee.

When the smartphone is within the relevant geo-fence, in optimal conditions, the platform collects GPS with 1 s resolution. Simultaneously, with the same resolution, the platform samples RSSI signal strength of BLE devices "visible" in the range of each smartphone.

We extracted the trajectories of both test passengers and buses between 1st April and 1st July. 134 users generated a total of 4,584,000 GPS observations; three buses, 1,162,000 GPS observations, for a total of approximately $940 \text{ h} \cdot \text{bus}$ operations (see Fig. 7).

From the remaining set of data we extracted the sub-set of observations containing at least one BLE observation, for a total of 195,000 GPS observations (see Fig. 6). This set present the maximum BLE resolution available, while the corresponding GPS resolution is below the maximum resolution available within the dataset. No labels are available for this set. Fig. 4 depicts the speed distribution of different transportation modes present in this subset. To highlight the differences in speed between different transport mode, we applied the exponential transformation. However, the black flat color shows that the speed distribution seems to be the same in all the cases, except for some cars (see black magnified detail).

Outside the passengers' set, we generated a set of records counting 59,000 observations which are part of a specific experiment where seven components of the project's staff collected via smartphone a high quality BIBO labels and observations set (see Fig. 5), following the same methodology of Shankari et al. for MobilityNet dataset collection [Shankari et al., 2020]. Thus, to avoid bias in the labels, we provided instructions on precise origin-destination sequences, divided in three different trip-groups. Each staff member has been randomly assigned to a trip-group. After watch synchronization, during the experiment, each staff member annotated the hour and minute each time s/he boarded or alighted a bus.

²Information provided to users before recruitement, access on 03-09-2021

³LINC DTU at GooglePLay, access on 03-09-2021

⁴LINC DTU at Appstore, access on 03-09-2021



Figure 4: Subset of GPS points presenting at least one BLE device reading; color map based on e^{speed} shows that buses and other modes in the area have the same speed distribution–i.e., walk and bike–few trajectories recorded from car are the only exception.



Figure 5: GPS points from smartphones, color map based on spatial density shows bus stops and bus deposit.



Figure 6: Subset of GPS points presenting at least one BLE device reading; points spatial distribution shows higher density at the bus stops, bus deposit and some buildings.



Figure 7: GPS points from buses, spatial distribution shows higher density at the bus stops, bus deposit.



Figure 8: Be-In (BI) clusters identified on smartphone data clustering CEMWA latent space with DBSCAN, and colored with ground truth labels. Red color depicts users inside the bus; blue color, users outside the bus.



Figure 9: Be-Out (BO) clusters identified on smartphone data clustering CEMWA latent space with DBSCAN, and colored with ground truth labels. Red color depicts users inside the bus; blue color, users outside the bus.

2.4.2 Experiment setup

Table 1 describes experimental setup for the evaluation of supervised baselines, for ablation study of various unsupervised architectures, and for the model we propose in this work. We applied a trajectory segmentation considering each pair of points beyond 120 s time-range, or where the space variation over time variation is beyond 120 m/s, the end of a segment and the beginning of the next segment. After segmentation, for each segment we applied a sliding window including 9 consecutive points and 1 step stride. CEMWA, MWA and WA process the resulting tensor straightly, using convolutions. Instead, Random Forest and XGboost require an intermediate process to extract traditional features from the 9 step windows contained in each segment, computed at each slide, applying the same stride of 1 step.

We setup the same conditions for both baselines and proposed methods. Comparing supervised and unsupervised classifiers in this setting is subject to the limitation of labeled dataset. As we want to provide performance distributions instead of points, with supervised methods we apply leave-one-out validation method, while with the unsupervised methods we apply a hold out method. In the first case we train the model with all the users belonging to the labeled observations except one, which represent the test set. In the test set we rotate all the users available. Thus, the main scores can presented as mean \pm standard deviation. In the second case, we train the model with the unlabeled observations, and without performing DBSCAN clustering. Then we use the model including DBSCAN to classify—off the sample—the labeled observations. Similarly, we can present the main scores as mean \pm standard deviation. Consequently, we can compare these scores even though the training process is quite different.

This setup assumes that the ground truth quality is stable and high. As we mentioned, the labels collection method we used can guarantee a higher quality level on the labels. Unlike the case where ground truth is collected from passengers, the project's staff followed instructions and was not subject to, e.g., recall bias, and less likely to suffer systematic and random distractions. Therefore, to provide an exhaustive picture for performance, we train these supervised methods adding some noise in the training set, i.e., flipping a controlled percentage of labels. We sample the number of errors per user from a Poisson distribution and we flip labels accordingly. The test set is not affected. Therefore, applying a Monte Carlo evaluation based on 100 loops per experiment, and on the same setup described in Table 1, we can estimate the sensitivity to labels noise. This problem does not affect the unsupervised methods, which use Bluetooth RSSI signal as pseudo-labels instead (see Table 1, Signals row).

3 Results and Discussion

After a manual optimization process of CEMWA, MWA, and WA, we yield optimal performance with the combination of hyperparameters described in Table 2. As opposed to CEMWA, MWA and WA converge to a relatively lower loss, and overfitting is higher. Although the three models have the same number of parameters, we record differing computation times for the training phase (which might be justified by concurrent processing on GPU). Compared to MWA and WA, CEMWA achieves substantially better scores, with higher mean and inferior standard deviation. (5) yields the results we present, while (6) seems not effective in this use case. We apply the same penalization across all three models during back-propagation to rebalance BI and BO classes when computing the WAE loss within the optimizer. Rather than the Precision score, the Recall score of the BI class seems to provide an essential contribution to the overall superior performance of CEMWA.

The supervised methods we evaluate are performing very well. XGboost presents a slightly higher score than CEMWA but with a slightly larger standard deviation. The two models seem to have comparable performance in terms of computation time. There seems to be the following differences. In optimal conditions and ground truth quality, XGboost appears to record a substantially higher precision score, but a lower recall score than CEMWA. Under the same conditions, Random Forest seems comparable with MWA and WA, or better. But we should not forget the impact of wrong labels in the training process of supervised methods such as XGboost and Random Forest. This problem does not affect unsupervised methods like CEMWA.

To test the sensitivity of XGboost and Random Forest to noise in the labels, we run a Monte Carlo evaluation. Results show that beyond 10% flipped labels during training leads to substantial performance degradation. This rapid degradation is of critical importance when labels are collected directly from passengers. Consequently, the trade-off between the cost and the quality of labels collection critically impacts the scalability potential of supervised methods. Figure 10 depicts the impact of wrong labels on the classifiers performance: When users provide wrong labels to less than 1 segment in average–where a segment is defined according to the GPS Trajectory Segmentation of Table 1–the performance of supervised classifiers drops dramatically compared to CEMWA.

This configuration provides potential for enhancing smartphone battery efficiency and user privacy, because: (i) Smartphones would listen to Bluetooth, while keeping GPS up, with minimum resolution, just enough to avoid GPS cold start; (ii) Bluetooth in proximity would trigger higher resolution GPS, only when necessary.



Figure 10: Impact of wrong labels on supervised classifiers training (F1 score macro average).



Figure 11: Impact of wrong labels on supervised classifiers training (F1 score weighted average).





1.75

1.50

0.25 0.50 0.75 1.00 1.25 mean number of segments with wrong labels during training

CEMWA
 CEMWA
 XGboost
 Random Forest
 Random

0.00

0.4

0.3

	Table 1: Experiment Setup					
	Supervised Baseline XG-Boost Random Forest	Unsupervised Baseline MWA (Fig. 2) WA (Fig. 3)	CEMWA (Fig. 1)			
Smartphone Set GPS + BLE Android + iOS	59,000 labelled observations 7 users	328,000 tot observations 59,000 labelled 134 tot users				
Buses set	1,162,000 observations, 940 h					
Signals	Speed, Longitude, Latitude, Timestamp from GPS	Speed, Longitude, Latitude, Timestamp from G RSSI and Timestamp from BLE devices				
Use of Ground Truth Labels	For training and evaluation	For evaluation only				
GPS Trajectory Segmentation	time gap between points >120 s deterpoints representing speed >45 m/s det	s determines a new segment 1/s determine a new segment				
Data Cleansing	Segments <10 consecutive point	s are discarded				
Observation Imputation	Imputation with Exponential Weighted Moving Average and Masking	g Masking Only				
Basic Feature Extraction	time-, space-gap, and bearing between each pair of GPS points, GPS di	stance between smartphone and buses within 1 s range				
Time Series Sliding Window	moving window of 9 consecutive steps se	steps segment, and 1 step stride				
Feature Extraction on Sliding Window	Mean value Max value Min value Position of the minimum value Position of the maximum value Amplitude between min and max value Ampler of points beyond one std dev. Number of points below one std dev. Number of points above one std dev. Number of peaks in the mvoing window Number of peaks half sliding window Number of peaks above 1 one std dev. Peak distance within sliding window Slope	$\label{eq:constraint} \begin{array}{c} \text{None.} \\ \text{ANN performs features extraction.} \\ \text{Encoder, 1 convolutional neural network} \\ \text{Decoder, 1 transposed convolutional neural network} \\ \text{Decoder, 1 transposed convolutional neural network} \\ \lambda \in [10^{-4}, 1] \\ \text{Batch Size: } \in [16, 1024] \\ \text{true sample size: } \in [10, 100] \\ \text{Learning Rate: } \in [10^{-5}, 10^{-1}] \\ \text{Epochs: } \in [10, 100] \end{array}$				
Performance Evaluation Method	Leave-one-out: One user in the test-set Training-set is the complementar set. Repeated rotating each user in test-set.	Hold-out: Training- and validation-set from unlabelled-set. Test-set corresponding to the labelled-set.				
Method performance distribution	Given by performance on individual users of whole the labelled set.					
Performance Metric	ic AUC ROC, F1-score, Precision, Recall, Accuracy					

Table 2: Encoder/Decoder CNN architecture hyperparameters, final configuration for CEMWA, EMWA, and WA.

Encoder	
Convolutional Neural Network (CNN) Layers	1
Activation Function	Rectified Linear Unit
Fully connected Layers	0
Dropout	0.25
Decoder	
Transposed CNN Layers	1
Activation Function	Leaky Rectified
	Linear Unit
Fully connected Layers	0
Dropout	0.25
Optimizer	Adam
Epochs	50
Batch Size	32
Learning Rate	10^{-4}
Dropout	0.25

In practice, after cause-effect training with encoder-decoder architecture and clustering-where GPS compression is trained reconstructing BLE and vice-versa-CEMWA could be deployed as follows. During operations, one CEMWA's encoder compresses GPS, while a separate encoder compress Bluetooth. The two independent compressed representation are joined into one. The proximity between the resulting representation and the clusters determine whether the observation belong to BI or BO class.

For applications where disruptions are unlikely-thus we expect a stable process in time-the amortization of high-quality ground truth could rely on a longer time horizon. An established metro line for example, is unlikely to experience changes frequently. In contrast, bus services are subject to continuous disruptions, e.g., roadworks and traffic congestion. Therefore, a supervised BIBO classifier could be a good choice in the first case. However, the unsupervised BIBO classifier seems better in the second case. Results rely mainly on the smartphone-bus-distance. This feature can be challenging to compute off-line, especially when a large number of passengers and vehicles are active. However, a federated-learning design [3rd Generation Partnership Project (3GPP), 2021] would solve the problem, and allow the computation of features online.

Assuming smartphones' future market penetration stable, and relying on adversarial sensors architectures, we show an approach to substitute manually collectible labels. This approach has vast potential; for example, BLE beacons contraposed to GPS within a CEMWA architecture would enable ticketless transit across any public transportation system, and large-scale deployment, even for applications subject to frequent disruptions. In addition to the beforementioned use case, we suggest road and bridge tolls or sharing mobility services like cars, bikes, or scooters. A BIBO system also supports visually impaired people to chose to board the right bus from the bus stop or to alight at the right stop from the bus. It could facilitate the integration across multiple service providers, operating mostly on software instead of physical infrastructure, even integrating with existing CICO and WIWO systems.

	Table	3: Results	with op	timal C	Ground	Truth fo	or method	evaluat	ion and	training	of supervis	ed algor	ithms
odel	Task	Labeled Observations	Unlabeled Observations	Precision	Recall	F1-s macro	core weighted	Accuracy	AUC ROC	Model Parameters	Ĉomputat Troining	ion-time Feature	Evaluati Methor

Model	Task	Observations	Observations	Frecision	Recail	macro average	weighted average	Accuracy	AUC KOC	Parameters	Training	Feature Extraction	Method
CEMWA	BI BO	13,154 45,327	191,556	0.77 0.97	0.89 0.92	0.88 ± 0.06	0.92 ± 0.04	0.91 ± 0.04	0.91	5028	97 min on 191,556 set		in Hold out (scrores distribution on labeled-set, comparable with
MWA	BI BO	13,154 45,327	191,556	0.52 0.89	0.66 0.82	0.72 ± 0.26	0.79 ± 0.22	0.79 ± 0.26	0.72	5028	47 min on 191,556 set	$< 1 \min$	
WA	BI BO	13,154 45,327	191,556	0.76 0.87	0.53 0.95	0.77 ± 0.15	0.85 ± 0.08	0.86 ± 0.08	0.74	4932	23 min on 191,556 set		leave-one-out)
XG-boost	BI BO	13,154 45,327	not applicable	0.84 0.93	0.78 0.95	0.90 ± 0.07	0.93 ± 0.06	0.93 ± 0.05	0.98	not	$< 1 \min$	31 min on	Leave-one-out
Random Forest	BI BO	13,154 45,327		0.90 0.85	0.43 0.99	0.82 ± 0.11	0.88 ± 0.07	0.90 ± 0.05	0.90	applicable		56,461 Set	
Random Classifier	BI BO	13,154 45,327		0.24 0.76	0.50 0.50	0.46 ± 0.01	0.54 ± 0.01	0.50 ± 0.003	0.50		not applicable	not applicable	not applicable

4 Conclusion

This paper focuses on an implicit tracking system to detect whether a passenger is inside or outside the transport network. To avoid using labels in the classifier training, we leverage a novel artificial neural network architecture learning the cause-effect relationship between two independent sensors measuring the same event. We call this approach CEMWA. In optimal conditions and with high-quality ground truth, CEMWA's performance is comparable or better than both supervised and unsupervised baselines. CEMWA and XGboost performance evaluated with optimal knwoledge on BIBO ground truth seem promising for public transport ticketing in general. In situations with noisy ground truth–such as transport services subject to disruption or surveys where passengers lack the ticket payment as an incentive to provide exact ground truth–we show that supervised classifiers' performance degrades. Supervised methods' tolerance to noisy labels is case specific. However, the issue does not affect CEMWA by design. Consequently, this unsupervised method is both scalable and fulfills the requirements for use-cases where, e.g., frequent service disruptions may lead to the need for regular labels' collection. Future research will investigate in few directions: (i) The extension of a sensor-to-sensor validation on new signals and neural network architectures, the sensitivity to labeling noise; (ii) The introduction of sensitivity to noise as a performance index to evaluate and compare supervised methods; and (iii) The connection between dry machine learning scores of our BIBO classifier and key performance index assessing automatic fare collection systems with BIBO.

Acknowledgment

This project is co-financed by the European Regional Development Fund through the Urban Innovative Actions Initiative.

References

- Wolfgang Narzt, Stefan Mayerhofer, Otto Weichselbaum, Stefan Haselbock, and Niklas Hofler. Be-in/be-out with bluetooth low energy: Implicit ticketing for public transportation systems. *Ieee Conference on Intelligent Transportation Systems, Proceedings, Itsc*, 2015-:7313345, 2015. ISSN 21530017, 21530009. doi:10.1109/ITSC.2015.253.
- Sampo Hietanen. Mobility as a service. the new transport model, 12(2):2–4, 2014.
- David A. Hensher and Corinne Mulley. Hensher, d.a. and mulley, c. mobility bundling and cultural tribalism might passenger mobility plans through maas remain niche or are they truly scalable? *Transport Policy*, 100:172–175, 2021. ISSN 0967-070X. doi:https://doi.org/10.1016/j.tranpol.2020.11.003. URL https://www.sciencedirect.com/science/article/pii/S0967070X20309203.
- Warwick Goodall, Tiffany Dovey, Justine Bornstein, and Brett Bonthron. The rise of mobility as a service. *Deloitte Rev*, 20:112–129, 2017.
- Matthias H. Wirtz and J. A.N. Klähr. Smartphone based in/out ticketing systems: A new generation of ticketing in public transport and its performance testing. *Wit Transactions on the Built Environment*, 182:351–359, 2019. ISSN 17464498, 17433509. doi:10.2495/UT180321.
- Valentino Servizi, Camara Francisco Pereira, Karen Marie Anderson, and Anker Otto Nielsen. Transport behaviormining from smartphones: a review. *European Transport Research Review*, 2021a. doi:10.1186/s12544-021-00516-z. URL https://doi.org/10.1186/s12544-021-00516-z.
- Oana Baescu and Hjalmar Christiansen. The Danish National Travel Survey Annual Statistical Report TU0619v2. DTU Management, 2020. doi:10.11581/dtu:00000034.
- Youjing Cui and Shuzhi Sam Ge. Autonomous vehicle positioning with gps in urban canyon environments. *Ieee Transac*tions on Robotics and Automation, 19(1):15–25, 2003. ISSN 2374958x, 1042296x. doi:10.1109/TRA.2002.807557.
- Lin Wang, Hristijan Gjoreski, Mathias Ciliberto, Sami Mekki, Stefan Valentin, and Daniel Roggen. Enabling reproducible research in sensor-based transportation mode recognition with the sussex-huawei dataset. *IEEE Access*, 2019. ISSN 21693536. doi:10.1109/ACCESS.2019.2890793.
- E. Foxlin. Inertial head-tracker sensor fusion by a complementary separate-bias kalman filter. In *Proceedings of the IEEE* 1996 Virtual Reality Annual International Symposium, pages 185–194, 1996. doi:10.1109/VRAIS.1996.490527.
- Andreas Bjerre-Nielsen, Kelton Minor, Piotr Sapiezynski, Sune Lehmann, and David Dreyer Lassen. Inferring transportation mode from smartphone sensors: Evaluating the potential of wi-fi and bluetooth. *Plos One*, 15(7), 2020. ISSN 19326203.
- Ivan Malmberg. An analysis of iBeacons and critical minimum distances in device placement. PhD thesis, 2014. URL http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-187925.
- Valentino Servizi, Dan Roland Persson, Per Bækgaard, Hannah Villadsen, Inon Peled, Jeppe Rich, Francisco C Pereira, and Otto A Nielsen. Context-aware sensing and implicit ground truth collection: Building a foundation for event triggered surveys on autonomous shuttles: Artikel. In *Proceedings from the Annual Transport Conference at Aalborg* University, volume 28, 2021b.
- Kalyanaraman Shankari, Jonathan Fuerst, Mauricio Fadel Argerich, Eleftherios Avramidis, and Jesse Zhang. Mobilitynet: Towards apublic dataset for multi-modal mobility research. *Climate Change AI*, 2020.
- Valentino Servizi, Niklas Christoffer Petersen, Francisco Camara Pereira, and Otto Anker Nielsen. Stop detection for smartphone-based travel surveys using geo-spatial context and artificial neural networks. *Transportation Research Part C: Emerging Technologies*, 121:102834, 12 2020. ISSN 0968090X. doi:10.1016/j.trc.2020.102834. URL https://linkinghub.elsevier.com/retrieve/pii/S0968090X20307385.
- Sina Dabiri and Kevin Heaslip. Inferring transportation modes from GPS trajectories using a convolutional neural network. *Transportation Research Part C: Emerging Technologies*, 86(November 2017):360–371, 2018. ISSN 0968090X. doi:10.1016/j.trc.2017.11.021. URL https://doi.org/10.1016/j.trc.2017.11.021.
- Maria Cornacchia, Koray Ozcan, Yu Zheng, and Senem Velipasalar. A survey on activity detection and classification using wearable sensors. *Ieee Sensors Journal*, 17(2):7742959, 2017. ISSN 15581748, 1530437x, 23799153. doi:10.1109/JSEN.2016.2628346.
- Anil NP Koushik, M. Manoj, and N. Nezamuddin. Machine learning applications in activity-travel behaviour research: a review. *Transport Reviews*, 0(0):1–24, 2020. doi:10.1080/01441647.2019.1704307. URL https://doi.org/10.1080/01441647.2019.1704307.
- Ivana Semanjski, Sidharta Gautama, Rein Ahas, and Frank Witlox. Spatial context mining approach for transport mode recognition from mobile sensed big data. *Computers, Environment and Urban Systems*, 66:38–52, 2017. ISSN 01989715. doi:10.1016/j.compenvurbsys.2017.07.004.

- Lijuan Zhang, Sagi Dalyot, Daniel Eggert, and Monika Sester. Multi-stage approach to travel-mode segmentation and classification of gps traces. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences:*[Geospatial Data Infrastructure: From Data Acquisition And Updating To Smarter Services] 38-4 (2011), Nr. W25, 38(W25):87–93, 2011.
- A. Kotanen, M. Hännikäinen, H. Leppäkoski, and T. D. Hämäläinen. Experiments on local positioning with bluetooth. Proceedings Itcc 2003, International Conference on Information Technology: Computers and Communications, page 1197544, 2003. doi:10.1109/ITCC.2003.1197544.
- Fazli Subhan, Halabi Hasbullah, Azat Rozyyev, and Sheikh Tahir Bakhsh. Indoor positioning in bluetooth networks using fingerprinting and lateration approach. 2011 International Conference on Information Science and Applications, Icisa 2011, page 5772436, 2011. doi:10.1109/ICISA.2011.5772436.
- Liang Chen, Heidi Kuusniemi, Yuwei Chen, Jingbin Liu, Ling Pei, Laura Ruotsalainen, and Ruizhi Chen. Constraint kalman filter for indoor bluetooth localization. 2015 23rd European Signal Processing Conference, Eusipco 2015, page 7362717, 2015. ISSN 20761465. doi:10.1109/EUSIPCO.2015.7362717.
- Fazli Subhan, Halabi Hasbullah, and Khalid Ashraf. Kalman filter-based hybrid indoor position estimation technique in bluetooth networks. *International Journal of Navigation and Observation*, 2013:570964, 2013. ISSN 16876008, 16875990. doi:10.1155/2013/570964.
- Héctor José Pérez Iglesias, Valentín Barral, and Carlos J. Escudero. Indoor person localization system through rssi bluetooth fingerprinting. 2012 19th International Conference on Systems, Signals and Image Processing, Iwssip 2012, page 6208163, 2012. ISSN 21578672.
- Ali Yassin, Youssef Nasser, Mariette Awad, Ahmed Al-Dubai, Ran Liu, Chau Yuen, Ronald Raulefs, and Elias Aboutanios. Recent advances in indoor localization: A survey on theoretical approaches and applications. *Ieee Communications Surveys and Tutorials*, 19(2):7762095, 2017. ISSN 1553877x, 2373745x. doi:10.1109/COMST.2016.2632427.
- Nadine Schuessler and Kay W. Axhausen. Processing raw data from global positioning systems without additional information. *Transportation Research Record*, 2105(1):28–36, 2009. doi:10.3141/2105-04. URL https://doi.org/10.3141/2105-04.
- Niklas Christoffer Petersen, Anders Parslov, and Filipe Rodrigues. Short-term bus travel time prediction for transfer synchronization with intelligent uncertainty handling. arXiv preprint arXiv:2104.06819, 2021.
- Richard Zhang, Phillip Isola, and Alexei A. Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction, 2016.
- Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.
- Arthur Gretton, Karsten Borgwardt, Malte J. Rasch, Bernhard Scholkopf, and Alexander J. Smola. A kernel method for the two-sample problem, 2008.
- Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein auto-encoders. *arXiv preprint* arXiv:1711.01558, 2017.
- Cédric Villani. Topics in optimal transportation. Number 58. American Mathematical Soc., 2003.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
- Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. arXiv preprint arXiv:1511.05644, 2015.
- Bharath K Sriperumbudur, Kenji Fukumizu, and Gert RG Lanckriet. Universality, characteristic kernels and rkhs embedding of measures. *Journal of Machine Learning Research*, 12(7), 2011.
- Aristidis Likas, Nikos Vlassis, and Jakob J. Verbeek. The global k-means clustering algorithm. *Pattern Recognition*, 36 (2):451-461, 2003. ISSN 0031-3203. doi:https://doi.org/10.1016/S0031-3203(02)00060-2. URL https://www.sciencedirect.com/science/article/pii/S0031320302000602. Biometrics.
- Kamran Khan, Saif Ur Rehman, Kamran Aziz, Simon Fong, and S. Sarasvady. Dbscan: Past, present and future. In *The Fifth International Conference on the Applications of Digital Information and Web Technologies (ICADIWT 2014)*, pages 232–238, 2014. doi:10.1109/ICADIWT.2014.6814687.
- Muhammad S. Osman, Adnan M. Abu-Mahfouz, and Philip R. Page. A survey on data imputation techniques: Water distribution system as a use case. *IEEE Access*, 6:63279–63291, 2018. doi:10.1109/ACCESS.2018.2877269.

Iman M. Almomani, Nour Y. Alkhalil, Enas M. Ahmad, and Rania M. Jodeh. Ubiquitous gps vehicle tracking and management system. In 2011 IEEE Jordan Conference on Applied Electrical Engineering and Computing Technologies (AEECT), pages 1–6, 2011. doi:10.1109/AEECT.2011.6132526.

3rd Generation Partnership Project (3GPP). Study on traffic characteristics and performance requirements for AI/ML model transfer, 22.874, 2021. URL https://portal.3gpp.org.