Coupled IGMM-GANs with Applications to Anomaly Detection in Human Mobility Data

DANIEL SMOLYAK, University of Maryland, USA
KATHRYN GRAY, University of Colorado, USA
SARKHAN BADIRLI and GEORGE MOHLER, Indiana University Purdue University Indianapolis, USA

Detecting anomalous activity in human mobility data has a number of applications including road hazard sensing, telematics based insurance, and fraud detection in taxi services and ride sharing. In this paper we address two challenges that arise in the study of anomalous human trajectories: 1) a lack of ground truth data on what defines an anomaly and 2) the dependence of existing methods on significant pre-processing and feature engineering. While generative adversarial networks seem like a natural fit for addressing these challenges, we find that existing GAN based anomaly detection algorithms perform poorly due to their inability to handle multimodal patterns. For this purpose we introduce an infinite Gaussian mixture model coupled with (bi-directional) generative adversarial networks, IGMM-GAN, that is able to generate synthetic, yet realistic, human mobility data and simultaneously facilitates multimodal anomaly detection. Through the estimation of a generative probability density on the space of human trajectories, we are able to generate realistic synthetic datasets that can be used to benchmark existing anomaly detection methods. The estimated multimodal density also allows for a natural definition of outlier that we use for detecting anomalous trajectories. We illustrate our methodology and its improvement over existing GAN anomaly detection on several human mobility datasets, along with MNIST.

CCS Concepts: • **Computing methodologies** → **Adversarial learning**; *Anomaly detection*; Neural networks; • **Applied computing** → *Transportation*;

Additional Key Words and Phrases: Bidirectional generative adversarial nets, unsupervised learning, infinite gaussian mixture models, mahalanobis distance

ACM Reference Format:
Daniel Smolyak, Kathryn Gray, Sarkhan Badirli, and George Mohler. 2020. Coupled IGMM-GANs with Applications to Anomaly Detection in Human Mobility Data. *ACM Trans. Spatial Algorithms Syst.* 1, 1, Article 1 (January 2020), 14 pages. https://doi.org/10.1145/3385809

1 INTRODUCTION

Human mobility data has become widely available through the proliferation of sensing devices and motivates a variety of machine learning tasks. One area of research focuses on activity detection and classification [Shoaib et al. 2015], where GPS or other sensor time series are used to classify trajectories (e.g. walking vs. biking). While in activity detection an exhaustive set of classes is known, the related task of anomaly detection in mobility data [Zhang et al.

Authors' addresses: Daniel Smolyak, University of Maryland, College Park, Maryland, USA, dsmolyak@umd.edu; Kathryn Gray, University of Colorado, Boulder, Colorado, USA, kathryn.gray@colorado.edu; Sarkhan Badirli; George Mohler, Indiana University Purdue University Indianapolis, Indianapolis, Indiana, USA.

2011] involves discovering samples that may belong to a new class or are outliers in the dataset. Finding anomalies in GPS trajectory data has a wide range of applications including determining anomalous traffic patterns [Pan et al. 2013], finding taxi drivers who commit fraud [Chen et al. 2013a], and driver finger printing and hidden driver detection in usage based insurance [Enev et al. 2016; Ramaiah et al. 2016].

In the majority of activity and anomaly detection studies, significant pre-processing and feature engineering are used prior to classification or trajectory similarity comparisons. While end-to-end deep learning has been recently applied to activity recognition [Ramaiah et al. 2016], to our knowledge no work to date has applied deep learning to anomaly detection in human mobility data. Furthermore, unlike activity recognition where ground-truth data is available, anomaly definition is often vague and subjective. With no ground truth datasets available, it is difficult to compare benchmark models available for detecting anomalies in trajectory and other sensor data.

To overcome these challenges, we propose using Generative Adversarial Networks (GAN) [Goodfellow et al. 2014] that can simultaneously generate realistic trajectory data as well as detect anomalies. In Figure 1, we provide an overview of our coupled IGMM-GAN model. We use a Bidirectional GAN (BiGAN) that learns an encoder in addition to a generator neural network for transforming trajectory data into a latent space where outliers may be detected. While BiGAN is the basis of our model, it only provides an embedding (but no anomaly score). Furthermore, previous BiGAN based anomaly detection methods [Schlegl et al. 2017; Zenati et al. 2018] have trouble detecting unseen classes when data is multi-modal, as we will show in our experiments section. For example, a unimodal model will have trouble detecting a hand written three digit when trained on sevens and eights, because the anomalous three is closer to an eight than an eight is to a seven [Gray et al. 2018]. While in our past work [Gray et al. 2018] we developed the IGMM-GAN framework for detecting anomalous images, here we show how this methodology can be adapted for generative modeling and anomaly detection in time-series data, with the specific application of mobility modeling in mind. For this purpose we propose a novel BiGAN model that can generate and encode both GPS routes and speed profiles simultaneously.

Unlike these previous methods, we use an Infinite Gaussian Mixture Model (IGMM) to detect anomalies in the latent space through a Mahalanobis distance based scoring. We will show in the following sections that coupling BiGAN with IGMM achieves much improved anomaly scores. Our contribution in this paper is two-fold: 1) we show how GANs can be applied to mobility data to generate realistic vehicle trajectories and velocity profiles and 2) we show how GAN based anomaly detection can be improved using IGMM. The IGMM component is important because it can handle multiple modes and covariance structures in the latent space, while simultaneously estimating the number of clusters.

The outline of the rest of the paper is as follows. In Section 2, we review related work on anomaly detection in trajectory data and generative adversarial networks. In Section 3, we provide details on the IGMM-GAN model. In Section 4, we present experimental results applying our model to MNIST and several human trajectory datasets. We compare Area Under the Curve (AUC) scores of the IGMM-GAN against several recently proposed GAN based anomaly detection algorithms and also provide a qualitative analysis of the generated synthetic trajectories and anomalies of the IGMM-GAN.

2 RELATED WORK

In this section we review previous literature on methods of detecting anomalies in GPS data and the use of GANs for synthetic data generation, embedding, and anomaly detection.

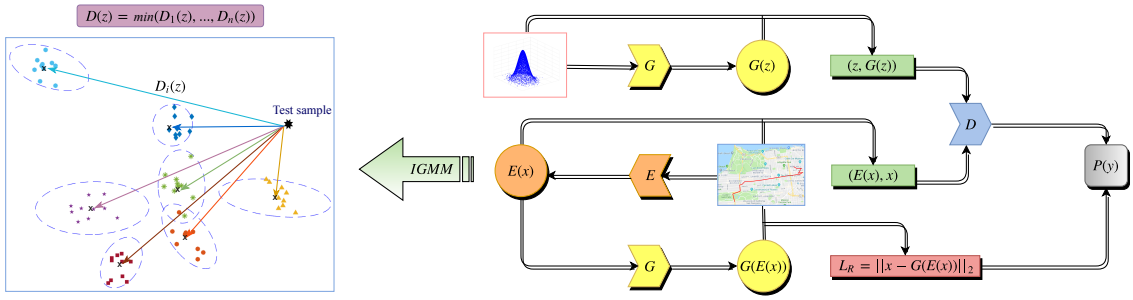

Fig. 1. The IGMM-GAN architecture. The generator network learns to transform Gaussian samples into synthetic trips, while the discriminator network learns to distinguish real from fake trips. Simultaneously, the encoder network learns the inverse mapping of the generator for trip embedding in the latent space. Finally, a multimodal Mahalanobis distance metric from the IGMM is used to detect outliers in unseen test data.

D	discriminator, classifies samples as real or fake
G	generator, maps Gaussian vector z to fake sample x
E	encodeer, maps real sample x to latex space variable z
D(z)	mahalanobis distance of a new sample to IGMM clusters
μ_i	mean of IGMM cluster i
Σ_i	covariance matrix of cluster i

Table 1. Variable definitions and descriptions for IGMM-GAN.

2.1 GPS Anomalies

Previous studies on anomaly detection in GPS trajectory data have focused both on the detection of city-wide traffic events and determining abnormal driving patterns on the level of an individual driver. On the more global level, [Donovan and Work 2015] used GPS data from New York City cab drivers to determine what effect natural disasters and other large scale disruptions had on the driving times in the city. In their work, traffic levels between regions are determined through the analysis of origin/destination pairs for taxi trips, allowing extrapolations to be made as to the traffic levels of the city as a whole. Combining GPS data and Twitter data, [Pan et al. 2013] were able to identify both the presence of increased traffic and its possible causes, such as accidents or weather. To find traffic anomalies, they analyzed road networks and driver trajectories as part of a graph, looking for time periods and routes where activity deviated from historical patterns. [Pang et al. 2013] also check for deviations in historical patterns, but instead use statistical spatio-temporal models of the number of taxis in a given region. Anomalies are found simply by looking for data points that are considered outliers within the statistical model of each region.

At the individual level, [Zhang et al. 2011] developed the Isolation Based Anomalous Trajectory (iBAT) detection method, which compares a driver's single trajectory to other trajectories with the same origin and destination. Dividing the space into a grid, iBAT calculates an anomaly score based on the similarity of the new route, in terms of the number of matching grid spaces. The related Isolation Based Anomalous Online Trajectory (iBOAT) detection method [Chen et al. 2013b], that consists of examining sub-trajectories through the lens of an adaptive working window, is able to complete similar calculations in real time and with more specificity, identifying precisely which part of a trajectory is anomalous. These two methods are able to detect not only increases in traffic, but also instances of fraud, where taxi drivers would deliberately take their passengers on longer routes in order to increase their earnings. [Bu et al.

2009] monitor for anomalies over continuous data streams as well, but instead use Euclidean distance to compare the trajectory under analysis to past and future trajectories. Much of their work focuses on using local clustering and other methods of optimization to complete these comparisons efficiently. Our goal in this paper is to develop an alternative, end-to-end learning framework, that does not require pre-processing trips on to grids or costly trip similarity scores. Instead, trips will be embedded in a Euclidiean space where a multimodal Mahalanobis distance is used to flag outliers.

2.2 GANs for synthetic data generation, embedding and anomaly detection

There have been several recent applications of GANs to synthetic data generation. [Molano-Mazon et al. 2018] applied GANs to generate and analyze data on the activity patterns of neurons, with the goal of creating data exhibiting the same statistics as the real data. To mitigate imbalanced data, [Salehinejad et al. 2017] used GANs to generate synthetic chest x-ray images. The synthetic data augmented samples of rare conditions and training with the generated data was shown to improve classifiers designed to identify these rare conditions.

[Gupta et al. 2018] generated movement trajectories based on socially acceptable behavior. These behaviors included passing and meeting people while walking and included elements such as speed and direction. GANs were then used to predict possible trajectories given a past trajectory. Latent representations of the trajectories were also used to determine direction and speed.

The closest past work on human mobility and deep data generation to ours is [Alzantot et al. 2017], in which a GAN-like architecture was used to generate synthetic acceleration data. Long-Short-Term-Memory networks coupled with mixture density networks were used to generate the acceleration time-series. However, the authors did not use a complete GAN architecture, in particular they did not use a discriminator to help train the generator and did not use an encoder for latent representations. The paper also focused solely on acceleration data, not GPS trajectories. In [Ouyang et al. 2018], the authors use GAN to generate synthetic mobility data on grids that preserve aggregate and individual statistics, but the goal is not anomaly detection and velocity profiles were not used. In [Gao et al. 2019], the authors use recurrent neural networks to embed trip data for the purpose of modeling the transitional distribution of points of interest and in [Zhou et al. 2019], a trajectory linking GAN is introduced to link trips together from the same user.

Several recent studies [Schlegl et al. 2017; Zenati et al. 2018] have successfully applied GANs for the purpose of anomaly detection making use of Bidirectional GAN (BiGAN) [Donahue et al. 2016]. These methods have fared favorably in anomaly detection compared to other deep embedding methods such as variational auto-encoders. However, we have found the existing GAN based anomaly detection methods (GANomaly and Efficient GAN Anomaly Detection) to have difficulties when the data is multimodal.

Several other recent studies have also made efforts at addressing this difficulty. [Ben-Yosef and Weinshall 2018] took the approach of modelling GANs from a Gaussian Mixture distribution as opposed to as single Gaussian, in order to better reflect the multiple classes of data that were being learned. This allowed them to generate a higher quality and more diverse set of synthetic images, and to specify the trade-off between quality and diversity. [Mukherjee et al. 2018] also improved on Vanilla GAN's representation of multimodal data by sampling from a more representative distribution for their synthetic data generation. By "sampling from a prior that consists of normal random variables cascaded with one-hot encoded vectors," they were able to better recover a smaller dimensional representation of various image data, and found that unsupervised clustering could be more effectively deployed here. While these papers both take a similar approach to our paper of being more cognizant of the multimodal nature of the data and taking advantage of the latent distribution, neither paper specifically applies their findings to anomalies or to human mobility data, as we will do in the following sections.

3 METHODOLOGY

In this section we provide background on generative adversarial networks and infinite Gaussian mixture models and then provide an overview of our proposed methodology.

3.1 GANs

GANs, first proposed in [Goodfellow et al. 2014], consist of a generator (G) network and a discriminator (D) network: the two follow the below minimax game, where the generator tries to minimize the $\log(1 - D(G(z)))$ term and the discriminator tries to maximize the $\log(D(x))$ term.

$$\max_D \min_G V(D,G) = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))]$$

The discriminator network improves the loss when it classifies a sample x correctly and $D(x)$ is the probability that x is real rather than generated data. Meanwhile, the generator network maps Gaussian noise z into synthetic data samples $G(z)$ (e.g. image or GPS trajectory). The generator attempts to minimize the discriminator loss by generating a fake sample $G(z)$ such that the discriminator labels the sample as real (hence the $1 - D(G(z))$ term).

3.2 BiGANs

Bidirectional GANs, first proposed by [Donahue et al. 2016], include an encoder (E) that learns the inverse of the generator. While the generator will learn a mapping from the latent dimension to data, the encoder will learn a mapping from data to the latent dimension. The discriminator then must classify pairs of the form $(G(z), z)$ or $(x, E(x))$ as real or synthetic, where z is noise from a standard distribution and x is real data.

$$\max_D \min_{G,E} V(D,G,E) = \mathbb{E}_{x\sim p_{data}(x)}\left[\mathbb{E}_{z\sim p_E(\cdot|x)}[\log D(x,z)]\right] + \mathbb{E}_{z\sim p_z(z)}\left[\mathbb{E}_{x\sim p_G(\cdot|z)}[\log(1 - D(x,z))]\right]$$

3.3 Anomaly Detection with BiGANs

As first proposed in the work by [Schlegl et al. 2017], and further developed by [Zenati et al. 2018], variants of GANs that also learn an inverse of the generator can be used to detect anomalous data. Specifically, after training a generator, discriminator, and encoder, an anomaly score can be calculated for each data sample, where a higher score indicates greater likelihood of belonging to the anomalous class. In the current state-of-the-art GAN based anomaly detection [Zenati et al. 2018], a combination of a reconstruction loss L_G and discriminator-based loss L_D is used to determine the anomaly score,

$$A(x) = \alpha L_G(x) + (1 - \alpha)L_D(x), \tag{1}$$

where the reconstruction loss is given by $L_G(x) = \|x - G(E(x))\|$ and the discriminator loss is given by the cross-entropy, $L_D(x) = \sigma(D(x, E(x)), 1)$. We refer to this algorithm as EGBAD (Efficient GAN based anomaly detection) and in [Zenati et al. 2018] the method is shown to outperform a variety of deep unsupervised models including anogan, variational auto-encoders, and deep auto-encoder GMM.

One main drawback of GAN based anomaly detection such as EGBAD and GANomaly is that their anomaly detection metric disregards the multimodal nature of the datasets (as we will show experimentally in the next section). Since real-world human mobility data contains many classes, drivers or pedestrians, with different characteristics and (assumption of) encoder from BIGAN preserves the clustering information in latent space, a single Gaussian model would be a very restrictive assumption for the data in latent space. Thus we relaxed this single Gaussian assumption by allowing the samples ($z' = E(x)$) in latent space to be sampled from a mixture of Gaussians, in particular IGMM. This in turn led the model capturing multiple modes with different Gaussian components and narrowing down the open space risk by dividing the entire space by these clusters (components). Thus the gap between normal and anomalous samples is further enlarged resulting remarkable detection performance boost. Outliers can then be detected using a Mahalanobis distance for the anomaly score. In particular, for a new data point x, $z' = E(x)$ in latent space, the Mahalanobis distance to cluster i (with mean μ_i and covariance Σ_i)) is given by,

$$D_M^i(z') = \sqrt{(z' - \mu_i)^T \Sigma_i^{-1} (z' - \mu_i)} \tag{2}$$

The anomaly score of sample x is then given by the minimum distance, $D(x) = \min_i D_M^i(z')$. As we will show in the next section, this Mahalanobis based anomaly score produces improved anomaly detection results over the cross-entropy loss in Equation 1.

3.4 Extending BiGAN for mobility data

Our model architecture is based on that of the BiGAN in [Donahue et al. 2016] and [Zenati et al. 2018] used for anomaly detection in image datasets. However, the architecture for the model needs to be customized for mobility time-series datasets and we detail our novel approach below.

The architecture for the model is given in Table 2. The encoder consists of an input layer taking in a $4 \times N$ time series of latitude, longitude, latitude velocity and longitude velocity (or an $N \times N$ image in the case of MNIST). For the mobility data we used time series of trip segments of length $N = 32$. Whereas the encoder consists of several convolution and dense layers, the generator makes use of convolution transpose layers to facilitate learning of the inverse of the encoder. The 2D convolution layers in the model are each followed by batch normalization and "Leaky ReLu" activation. The discriminator is slightly more complex, beginning as two separate models, one composed similarly to the encoder which takes the real and generated data as input, and one containing dense layers which takes the latent representation as input. These two networks are then concatenated, ending in two final dense layers and a sigmoid activation.

Furthermore, combining the ideas from [Akcay et al. 2018] and [Donahue et al. 2016], we add onto the existing architecture a reconstruction loss term, taking into account the ability of the encoder and generator to reproduce a real image. This loss term helps ensure that not only can the generator's images fool the discriminator, but also that the encoder and generator function as closely as possible to inverses of one another. This loss is defined as:

$$L_R = ||x - G(E(x))||_2$$

We use an Adam optimizer [Kingma and Ba 2014] with a learning rate of $lr = 1e^{-5}$ and $\beta = 0.5$. These parameters are sufficient for the generator and discriminator loss for our model to converge, similarly to the other models.

Layer	Units	BN	Activation	Kernel
$E(x)$				
Dense	768		ReLU	
Convolution	32	✓	ReLU	3×2
Convolution	64	✓	ReLU	3×2
Convolution	128	✓	ReLU	3×2
Dense	100			
$G(z)$				
Conv. Transpose	128	✓	ReLU	3×2
Conv. Transpose	64	✓	ReLU	3×2
Conv. Transpose	32			3×2
Dense	1		Linear	
$D(x)$				
Convolution	64		Leaky ReLU	3×2
Convolution	64	✓	Leaky ReLU	3×1
$D(z)$				
Dense	512		Leaky ReLU	
Concatenate				
$D(x, z)$				
Dense	1		Leaky ReLU	

Table 2. The architecture for our model, layer by layer. Units refer to number of filters in the case of convolution layers, and BN is Batch Normalization abbreviated.

3.5 Infinite Gaussian Mixture Model

Because our goal is end-to-end learning for anomaly detection, we use an infinite Gaussian mixture model (IGMM) [Rasmussen 1999] to automatically learn the number of clusters as well as the cluster means and covariances (μ_i, Σ_i) in the latent space for the anomaly score defined by Equation 2.

IGMM [Rasmussen 1999] is a Dirichlet Process Mixture Model in which the number of components can grow arbitrarily as data allows, hence the name Infinite Gaussian Mixture Model. IGMM assumes each cluster is modeled by a single Gaussian component and the base Dirichlet distribution serves as a prior for the parameters of these components (cluster mean μ and cluster covariance Σ). As the name Gaussian mixture suggests, the bi-variate prior, H, involves a Gaussian prior over mean vectors and Inverse-Wishart over covariance matrices. More precisely H can be written as follows,

$$H = N(\mu|\mu_0, \Sigma_0\kappa_0^{-1})W^{-1}(\Sigma|\Sigma_0, m) \tag{3}$$

where μ_0 is the mean of Gaussian prior, κ_0 is scaling constant that adjusts the dispersion of cluster center and parameter m dictates the expected shapes of clusters. Note that Normal and Inverse-Wishart distributions are conjugate, thus the posterior predictive distribution can be analytically derived, in the form of a multivariate Student-t distribution, by integrating out the component parameters $\{\mu_i, \Sigma_i\}$. For inference we utilize Collapsed Gibbs Sampling [Rasmussen 1999] due to the conjugacy between the model (Gaussian) and the prior (NIW). The generative model is illustrated in (4)

$$\begin{aligned} z_i &\sim N(z_i|\mu_i, \Sigma_i) \\ \{\mu_i, \Sigma_i\} &\sim G \\ G &\sim DP(\alpha H) \end{aligned} \tag{4}$$

where H is defined by Equation (3), z_i is the data point from cluster i and α is the concentration parameter of the Dirichlet process.

We provide an overview of our entire modeling process in Algorithm 1.

Algorithm 1: Coupled IGMM-GAN anomaly detection

Input: Training data $\mathcal{D}_{tr}$
Output: Anomaly score of each test sample
Transforming raw trips into structured data points in latent space
1. Train Bidirectional GAN from Figure 1 on $\mathcal{D}_{tr}$ to learn encoder E
2. Calculate $z = E(x)$ for each sample $x \in \mathcal{D}_{tr}$ to represent training trips in the latent space
Deriving class sufficient statistics of the training data from IGMM clusters
3. Run IGMM on these points in the latent space, z's, to cluster the training data: assume IGMM finds K clusters
4. Calculate sample mean μ_i and and covariance Σ_i for each $i \in \{1, 2, ..., K\}$
Anomaly score calculation
5. For a test sample $x' \in \mathcal{D}_{ts}$ calculate $z' = E(x')$ to get the latent dimension representation
6. Calculate Mahalanobis distance between z' and clusters from IGMM, $D_M^i(z')$, for all $i \in 1, 2, ..., K$ (Eq. 2)
7. Finally anomaly score of test sample x' is the min of these distances: $D_M(x') = \min_i D_M^i(z')$

4 EXPERIMENTAL RESULTS

In this section we describe our datasets in sub-section 4.1, the architecture of our model in sub-section 4.2, the hyper-parameters of our model in sub-section 4.3, and then the results of our experiments in subsections 4.4 to 4.7.

4.1 Datasets

In this section we run several experiments in order to compare our IGMM-GAN to several benchmark GAN based anomaly detection algorithms, as well as to evaluate the ability of GANs to generate realistic GPS trajectory data. Our experiments are run on the following datasets.

4.1.1 GeoLife GPS Trajectories. The Geolife GPS trajectory dataset [Zheng et al. 2008, 2010, 2009], consists of data collected from 180 users going about their daily routine movements. The data consisted of time-series in the form of latitude, longitude, and velocity sampled every few seconds from each user. A subset of the trajectory data had information on the mode of transportation, and for our experiments we used only those trajectories that could be confirmed as “car”. To create each data point, or "route", we took a sliding window of the data, with no overlap, so each route had the measured location of 24 consecutive measurements. This method of data segmentation is used in other similar studies of sequential data, such as [Ramaiah et al. 2016]. As in other GAN based anomaly detection studies, we use a held-out class (specific human subject) during training to define an anomalous class in testing (and then evaluate over different choices of the subject held out).

4.1.2 San Francisco Cabspotting. Similar to Geolife, the SF Cabspotting data set consists of GPS trajectories from different cab drivers in San Francisco [Piorkowski et al. 2009]. This data consisted of latitudes, longitudes, and times, which were then expanded to include latitude, longitude, latitudinal velocity, and longitudinal velocity. Each route was constructed with a sliding window, same as the Geolife data. Additionally, since the cab data had many trips that were not possible (speeds too fast for a normal vehicle, or locations that were not possible, such as the middle of the ocean),

these points were removed before training the model. Anomalous routes were found by training on a subset and finding the top anomaly scores.

4.1.3 MNIST. Previous GAN based anomaly detection studies have used MNIST (a dataset of handwritten numbers) [LeCun et al. 2010] for bench-marking competing methods. Anomalies are defined by leaving out a digit from training and assessing the AUC (or other classification metric) of the anomaly score on a test data set which includes the held-out digit.

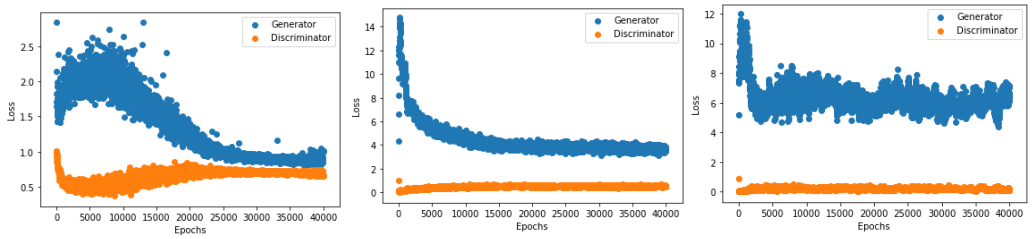

Fig. 2. Generator and discriminator loss by epoch for MNIST, digit 9 (Order: Ganomaly, EGBAD, Our Method)

4.2 Hyper-parameter tuning for IGMM

The hyper-parameters of IGMM are coarsely tuned to maximize the clustering macro-f1 score. As the data is not well balanced, macro-f1 was chosen to suppress the dominance of large classes. IGMM has 4 hyper-parameters, $\{\kappa_0, m, \mu_0, \Sigma_0\}$ to be tuned. To simplify the tuning process, the prior mean, μ_0, is set to the mean of data and we set Σ_0 to an identity matrix scaled by a parameter s. This left us with 3 parameters, $\{\kappa_0, m, s\}$ to tune. Parameter ranges and best triples are illustrated in Table 3. The number of sweeps in the inference is fixed at 500, with 300 used for the burn-in period. Label samples are collected in every 50 iteration after burn-in and aligned by the Hungarian method to render final cluster labels.

HP	Range
κ_0	$0.01; 0.1; 1; 10; 100$
m	$d+10; d+15; d+20; 5d; 10d; 100d$
s	$1; 3; 5; 7; 9$

(a) Parameter ranges used in tuning

HP	MNIST	Geolife
κ_0	0.1	0.1
m	$d+20$	$d+15$
s	7	5

(b) Best triples from tuning

Table 3. Ranges for tuning and best triples used in experiments. HP stands for hyper-parameters

We restricted created clusters to ones with more than 50 points as IGMM may generate artificial small clusters to fit in distribution.

4.3 Determining Anomaly Scores

Anomaly scores were determined by using IGMM on the encoded training data to determine the cluster means and covariance matrices. From there, an anomaly score was determined by the Mahalanobis distance to the nearest cluster. Figure 3 shows an example TSNE visualization [Maaten and Hinton 2008] of the GeoLife data in latent space colored by the anomaly scores. In this case, the driver that was not included in the training data is a clear outlying cluster reflected

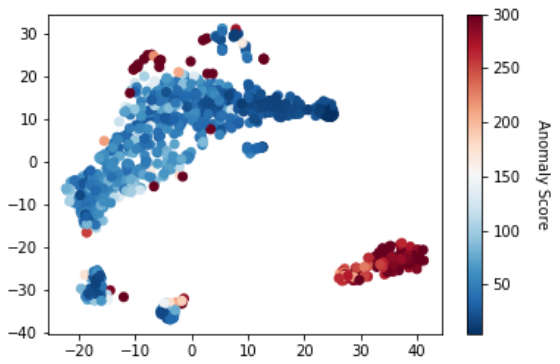

Fig. 3. TSNE visualization of the latent dimension with anomaly scores from Geolife data

by the anomaly scores. In the Geolife experimental section below we provide a detailed analysis of the performance of IGMM-GAN for detecting hidden drivers against several benchmark models.

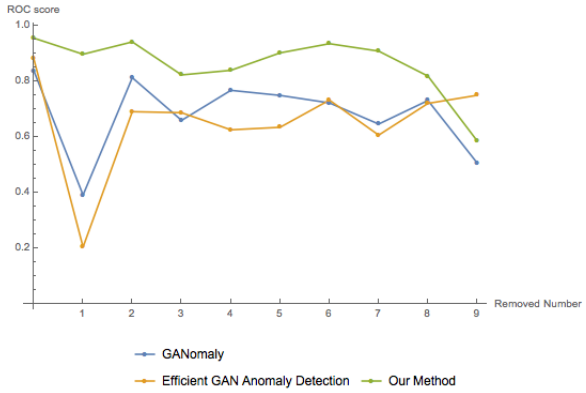

Fig. 4. ROC AUC scores with MNIST data

4.4 Improving MNIST Benchmarks with IGMM

Following the approach of [Zenati et al. 2018], we start by designating one digit as an anomalous class and remove it from the training dataset. For the remaining data we perform an 80/20 split into training and test sets and train the models for 40,000 epochs (where each epoch involves training on a random batch of 128 images). We repeat this process for each digit and for each anomaly detection method, scoring each method on its ROC AUC score. In Figure 4 we display the AUC scores of IGMM-GAN against GANomaly and EGBAD for each digit held out of testing. The IGMM-GAN significantly improves the AUC scores for the majority of digits held out, especially for digits 1 and 7. We note that for the uni-modal methods the AUC score is below 0.5 when the hand-written digit "1" class is held-out as the outlier. The overall variation in the latent space (as measured by the spectral radius of the covariance matrix) of the combined 1 & 7 class is actually smaller than the variation of several of the other single classes (for example 3, 6 and 8).

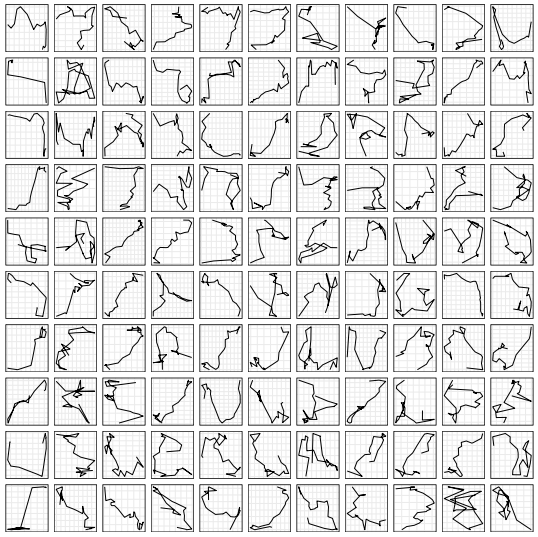

(a) Real and generated cab spotting routes.

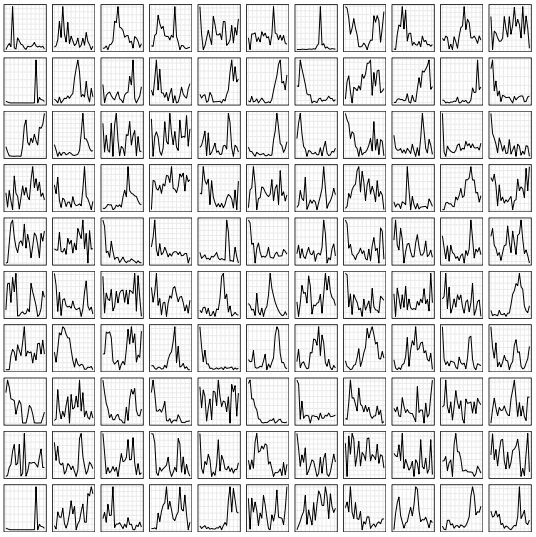

(b) Real and generated cab spotting speed profiles.

Fig. 5. In both figures, column 1 represents real values and columns 2-11 display 10 closest generated values (by Euclidean distance in the latent dimension).

So in this case it is possible for the AUC to go below .5. This is an example where a multi-modal outlier score is much preferred to a uni-modal score.

4.5 San Francisco Cabspotting Experiment

4.5.1 Generating Routes. We first examine the IGMM-GAN's ability to generate realistic synthetic trajectories. Note that the routes were created using a generator trained on both latitude and longitude as well as latitudinal and longitudinal velocity. Including the velocities gave slightly better results, as the sampling interval in the data was not uniform, hence with the added velocities, the generator could infer the time step.

We display 10 example real trajectories (Figure 5a) and speed profiles (Figure 5b) in the left most column of each grid along with the 10 closest trajectories and speed profiles (closest in latent space) to each real route. The generated trajectory data in Figure 5a matches the real data both upon visual inspection and in terms of convergence of the discriminator and generator loss. Qualitatively, both the synthetic and real routes, for latitude and longitude, are characterized by long stretches of straight lines or slight curves followed by segments of turns and short paths, as seen in Figure 5a. The velocities in the generated routes also follow a natural pattern of increase and decrease, as seen in Figure 5b. The generated velocities also follow general traffic patterns with longer distances between points often having higher velocity and slower velocities often being placed near the center of the city.

4.5.2 Finding Anomalous Routes. Because the cabspotting data consists of taxi drivers in San Francisco, who all drive similar routes all over the city, driver detection is more difficult with this dataset than with the geolife data analyzed in the next section. Therefore we used the latter for the quantitative analysis of anomalies. However, we are able to perform a qualitative analysis of the anomalous trips in the cabspotting data.

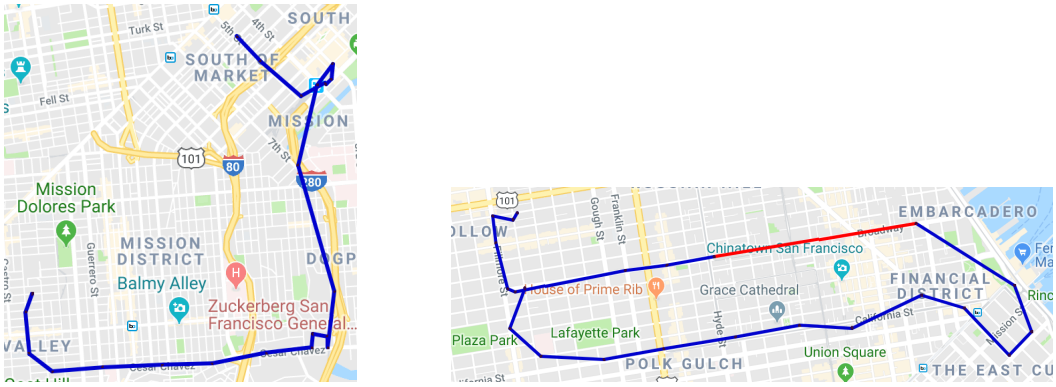

(a) Route with an unnecessary detour.

(b) Route with an unusual high speed in an urban area.

Fig. 6. Anomalous routes from Cabspotting dataset. The red line in figure (b) represents the route with speed over 60 mph.

Using the IGMM-GAN method, we assigned an anomaly score to each route in the cabspotting dataset. The first category of anomalous routes were those with GPS noise: in the dataset certain trips contain inaccurate points where the coordinates may have noise on the order of miles. Using both latitude, longitude and the respective velocities, these routes gave the highest anomaly scores. We therefore removed these trips prior to training our model. With the GPS noise trajectories removed, different types of anomalies emerged. The anomalies found were generally in two categories, anomalies in the route taken and anomalies in velocity.

Figure 6a is an example of the route anomaly. Note that in the bottom right corner the taxi takes a detour from the usual route. Other anomalies in the route taken included routes that had a large number of turns or ones that were right next to major roads, but did not use them.

In Figure 6b, the red part shows where the speed, instead of a normal speed for the area, was much higher (closer to highway speeds). Anomalies in the velocity generally were either speeds that were too fast for the area, such as in the image above, or speeds that were too slow for the area, such as driving below 30mph on a highway. Sometimes velocity anomalies were those where the acceleration suddenly changed and one segment would be slow and the next much faster. Thus the IGMM-GAN is able to detect not only uncommon routes, but also abnormal driver behaviors with the same model (using an end-to-end learning framework).

4.6 Geolife Hidden Driver Experiment

Our next experiment, with the GeoLife data, focuses on distinguishing the emergence of a new driver not contained in the training dataset. Here a similar method to the MNIST experiment was followed: one driver was defined as an anomaly and left out of training data, then included in the testing data as a separate class. In Figure 7, we compare the three methods from before on each of nine held out drivers. Here we see that GANomaly and EGBAD perform better on some drivers and worse on others compared to each other. However, IGMM-GAN is consistently at least as accurate as both methods and in 6 cases it shows the best AUC scores.

5 CONCLUSION

We have made three main contributions in this paper:

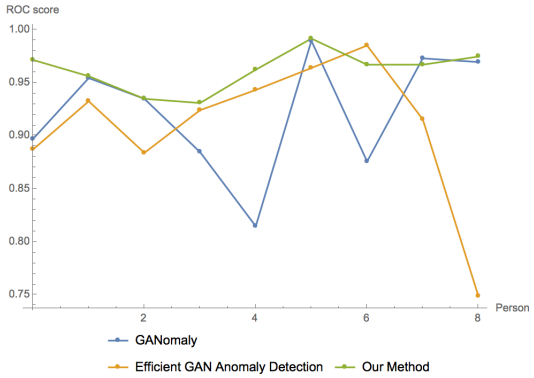

Fig. 7. ROC AUC Scores for Different Drivers from GeoLife Data

- We improved GAN based anomaly detection through the introduction of the multimodal IGMM-GAN.
- We showed that GAN can be successfully applied to generating synthetic human trajectories. Synthetic data sets of this type will be useful as ground truth to benchmark existing anomaly detection methods for human mobility and can also aid in data augmentation.
- We showed that the IGMM-GAN can find route-based anomalies, anomalous driver behavior, and detect hidden drivers.

We believe that the IGMM-GAN will serve as a complimentary tool to existing algorithms for anomaly detection in human mobility data that require spatial grids and feature engineering. Our method may also find application to anomaly detection in other domains where the data is multimodal.

In the present study there are also several limitations that may be the focus of future research. Here we focused on anomaly detection with respect to entire and individual trips. However, there are situations where anomaly detection may be desired on the fly within a trip. For example, to detect that a car is stolen, a sequentially updating model of how likely a car is driven by it's owner would be desirable. Also, anomalies may not only be specific to a particular vehicle, but could reflect hazards or changes to the environment or road network. Adapting the IGMM-GAN framework to detect these types of anomalies will be the focus of future research.

6 ACKNOWLEDGEMENTS

This work was supported in part by NSF grants ATD-1737996, REU-1343123, and SCC-1737585.

REFERENCES

Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. 2018. GANomaly: Semi-Supervised Anomaly Detection via Adversarial Training. *arXiv preprint arXiv:1805.06725* (2018).

Moustafa Alzantot, Supriyo Chakraborty, and Mani B. Srivastava. 2017. SenseGen: A Deep Learning Architecture for Synthetic Sensor Data Generation. *CoRR* abs/1701.08886 (2017). arXiv:1701.08886 http://arxiv.org/abs/1701.08886

Matan Ben-Yosef and Daphna Weinshall. 2018. Gaussian Mixture Generative Adversarial Networks for Diverse Datasets, and the Unsupervised Clustering of Images. *CoRR* abs/1808.10356 (2018). arXiv:1808.10356 http://arxiv.org/abs/1808.10356

Yingyi Bu, Lei Chen, Ada Wai-Chee Fu, and Dawei Liu. 2009. Efficient anomaly monitoring over moving object trajectory streams. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 159–168.

C. Chen, D. Zhang, P. S. Castro, N. Li, L. Sun, S. Li, and Z. Wang. 2013a. iBOAT: Isolation-Based Online Anomalous Trajectory Detection. *IEEE Transactions on Intelligent Transportation Systems* 14, 2 (June 2013), 806–818. https://doi.org/10.1109/TITS.2013.2238531

Chao Chen, Daqing Zhang, Pablo Samuel Castro, Nan Li, Lin Sun, Shijian Li, and Zonghui Wang. 2013b. iBOAT: Isolation-based online anomalous trajectory detection. *IEEE Transactions on Intelligent Transportation Systems* 14, 2 (2013), 806–818.

Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. 2016. Adversarial feature learning. *arXiv preprint arXiv:1605.09782* (2016).

Brian Donovan and Daniel B Work. 2015. Using coarse GPS data to quantify city-scale transportation system resilience to extreme events. *arXiv preprint arXiv:1507.06011* (2015).

Miro Enev, Alex Takakuwa, Karl Koscher, and Tadayoshi Kohno. 2016. Automobile driver fingerprinting. *Proceedings on Privacy Enhancing Technologies* 2016, 1 (2016), 34–50.

Qiang Gao, Goce Trajcevski, Fan Zhou, Kunpeng Zhang, Ting Zhong, and Fengli Zhang. 2019. DeepTrip: Adversarially Understanding Human Mobility for Trip Recommendation. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 444–447.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.

Kathryn Gray, Daniel Smolyak, Sarkhan Badirli, and George Mohler. 2018. Coupled IGMM-GANs for improved generative adversarial anomaly detection. In *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, 2538–2541.

Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. 2018. Social GAN: Socially Acceptable Trajectories with Generative Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

Yann LeCun, Corinna Cortes, and Christopher JC Burges. 2010. MNIST handwritten digit database. 2010. *URL http://yann. lecun. com/exdb/mnist* (2010).

Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.

Manuel Molano-Mazon, Arno Onken, Eugenio Piasini, and Stefano Panzeri. 2018. Synthesizing realistic neural population activity patterns using Generative Adversarial Networks. In *International Conference on Learning Representations 2018*.

Sudipto Mukherjee, Himanshu Asnani, Eugene Lin, and Sreeram Kannan. 2018. ClusterGAN : Latent Space Clustering in Generative Adversarial Networks. *CoRR* abs/1809.03627 (2018). arXiv:1809.03627 http://arxiv.org/abs/1809.03627

Kun Ouyang, Reza Shokri, David S Rosenblum, and Wenzhuo Yang. 2018. A Non-Parametric Generative Model for Human Trajectories.. In *IJCAI*. 3812–3817.

Bei Pan, Yu Zheng, David Wilkie, and Cyrus Shahabi. 2013. Crowd sensing of traffic anomalies based on human mobility and social media. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 344–353.

Linsey Xiaolin Pang, Sanjay Chawla, Wei Liu, and Yu Zheng. 2013. On detection of emerging anomalous traffic patterns using GPS data. *Data & Knowledge Engineering* 87 (2013), 357–373.

Michal Piorkowski, Natasa Sarafijanovoc-Djukic, and Matthias Grossglauser. 2009. A Parsimonious Model of Mobile Partitioned Networks with Clustering. In *The First International Conference on COMmunication Systems and NETworkS (COMSNETS)*. http://www.comsnets.org

Chetan Ramaiah, Allen Tran, Evan Cox, and George Mohler. 2016. Deep learning for driving detection on mobile phones. In *KDD Workshop on Machine Learning for Large Scale Transportation Systems*.

Carl Edward Rasmussen. 1999. The infinite gaussian mixture model. In *Advances in neural information processing systems*. 554–560.

Hojjat Salehinejad, Shahrokh Valaee, Tim Dowdell, Errol Colak, and Joseph Barfett. 2017. Generalization of deep neural networks for chest pathology classification in x-rays using generative adversarial networks. *arXiv preprint arXiv:1712.01636* (2017).

Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. 2017. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*. Springer, 146–157.

Muhammad Shoaib, Stephan Bosch, Ozlem Durmaz Incel, Hans Scholten, and Paul JM Havinga. 2015. A survey of online activity recognition using mobile phones. *Sensors* 15, 1 (2015), 2059–2085.

Houssam Zenati, Chuan Sheng Foo, Bruno Lecouat, Gaurav Manek, and Vijay Ramaseshan Chandrasekhar. 2018. Efficient GAN-based anomaly detection. *arXiv preprint arXiv:1802.06222* (2018).

Daqing Zhang, Nan Li, Zhi-Hua Zhou, Chao Chen, Lin Sun, and Shijian Li. 2011. iBAT: detecting anomalous taxi trajectories from GPS traces. In *Proceedings of the 13th international conference on Ubiquitous computing*. ACM, 99–108.

Yu Zheng, Quannan Li, Yukun Chen, Xing Xie, and Wei-Ying Ma. 2008. Understanding mobility based on GPS data. In *Proceedings of the 10th international conference on Ubiquitous computing*. ACM, 312–321.

Yu Zheng, Xing Xie, and Wei-Ying Ma. 2010. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data Eng. Bull.* 33, 2 (2010), 32–39.

Yu Zheng, Lizhu Zhang, Xing Xie, and Wei-Ying Ma. 2009. Mining interesting locations and travel sequences from GPS trajectories. In *Proceedings of the 18th international conference on World wide web*. ACM, 791–800.

Fan Zhou, Ruiyang Yin, Goce Trajcevski, Kunpeng Zhang, Jin Wu, and Ashfaq Khokhar. 2019. Improving human mobility identification with trajectory augmentation. *GeoInformatica* (2019), 1–31.