

Spatiotemporal Modeling and Forecasting at Scale with Dynamic Generalized Linear Models

Pranay Pherwani*
STR
Woburn, MA, USA
pranay.pherwani@str.us

Nicholas Hass*
STR
Woburn, MA, USA
nicholas.hass@str.us

Anna K. Yanchenko
STR
Woburn, MA, USA
anna.yanchenko@str.us

Abstract

Spatiotemporal data consisting of timestamps, GPS coordinates, and IDs occurs in many settings. Modeling approaches for this type of data must address challenges in terms of sensor noise, uneven sampling rates, and non-persistent IDs. In this work, we characterize and forecast human mobility at scale with dynamic generalized linear models (DGLMs). We represent mobility data as occupancy counts of spatial cells over time and use DGLMs to model the occupancy counts for each spatial cell in an area of interest. DGLMs are flexible to varying numbers of occupancy counts across spatial cells, are dynamic, and easily incorporate daily and weekly seasonality in the aggregate-level behavior. Our overall approach is robust to various types of noise and scales linearly in the number of spatial cells, time bins, and agents. Our results show that DGLMs provide accurate occupancy count forecasts over a variety of spatial resolutions and forecast horizons. We also present scaling results for spatiotemporal data consisting of hundreds of millions of observations. Our approach is flexible to support several downstream applications, including characterizing human mobility, forecasting occupancy counts, and anomaly detection for aggregate-level behaviors.

CCS Concepts

• **Mathematics of computing** → **Time series analysis; Bayesian computation.**

Keywords

spatiotemporal modeling, Bayesian state space models, human mobility data, modeling at scale, probabilistic forecasting

ACM Reference Format:

Pranay Pherwani, Nicholas Hass, and Anna K. Yanchenko. 2024. Spatiotemporal Modeling and Forecasting at Scale with Dynamic Generalized Linear Models. In *1st ACM SIGSPATIAL International Workshop on Geospatial Anomaly Detection (GeoAnomalies'24)*, October 29, 2024, Atlanta, GA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3681765.3698449>

*Both authors contributed equally to this research.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. Request permissions from owner/author(s).

GeoAnomalies'24, October 29, 2024, Atlanta, GA, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1144-2/24/10
<https://doi.org/10.1145/3681765.3698449>

1 Introduction

Spatiotemporal data occurs in a variety of different settings such as traffic management [25, 32], ecology [14], environmental monitoring [1] and epidemic control [7]. There are several characteristics of GPS trajectory spatiotemporal data that make accurate modeling of this data challenging, including sensor noise and uneven sampling rates of observations. Additionally, spatiotemporal data often exhibits heterogeneous dependencies over space, time, and between agents or entities in the trajectory data, and accurately capturing these dependencies while modeling spatiotemporal data at scale is a significant challenge. In this work, we propose an efficient, flexible, and scalable approach to model and forecast spatiotemporal data at scale using dynamic generalized linear models (DGLMs) [28, 29].

Our approach focuses on modeling human mobility data specifically and we are interested in modeling mobility data for hundreds of thousands of agents for months of time over areas of thousands of km². Our setting requires models that accurately capture normal human mobility to support a variety of downstream applications, including forecasting and anomaly detection. We require modeling approaches that are efficient, probabilistic, robust to various types of noise, and can incorporate important dependencies exhibited in the data, such as daily seasonality in human mobility. We choose DGLMs in this setting due to their ability to meet these requirements.

1.1 Related Work

Accurate models for human mobility data support many downstream applications, including traffic flow forecasting [2, 22, 25], urban planning [21], point of interest demand analysis [17], epidemic modeling [21], and anomaly detection. Related work in mobility modeling can largely be separated into modeling fine-grained agent movement or larger scale, population-level movement [2] and can also be separated by which modes of transportation are modeled, e.g. pedestrian movement, vehicular movement, etc. [21].

There are a variety of approaches to modeling spatiotemporal data generally and human mobility data specifically, ranging from traditional point process models to more recent deep learning approaches. Mechanistic models, such as the gravity model, model mobility flow between locations as inversely related to the distance between those locations [33]. Bayesian latent factor models were used to model demand patterns based on points of interest in [17]. Gaussian processes are commonly used statistical models for spatiotemporal data; recent extensions include capturing social-network dependencies between agents in animal movement data [20] and combining Gaussian processes with deep learning to predict origin-destination flows in taxi data [23]. Finally, deep learning approaches are becoming more popular for spatiotemporal data,

for example, to model and predict city-wide traffic patterns with LSTMs [22].

Dynamic generalized linear models are Bayesian state space models that are flexible to various observation distributions [28, 29]. They have been successfully applied in a variety of settings involving large-scale data, including retail forecasting [3, 30] and online advertising for modeling traffic flows through a website [5]. Critically, in our setting observations are preprocessed into occupancy counts in space-time bins and DGLMs can flexibly model count time series data with a variety of sparsity and over-dispersion levels [3]. Related work using state space models for spatiotemporal data includes traffic flow prediction and modeling with Kalman filters [10, 18], though only for settings that admit normally distributed observations, modeling rainfall data and water temperature [24] and modeling tornado reports and mortality ratios [11].

In this work, we focus on aggregate-level modeling of human mobility over long and short time horizons and our approach is applied to data with walking and driving movement patterns. While our results focus on multi-step ahead forecasting of occupancy counts, we propose a modeling approach that is flexible to a variety of downstream applications. To this end, we require models that are probabilistic to accurately quantify uncertainty, can be used for multi-step ahead forecasting, are able to flexibly incorporate external information (e.g. point of interest data, holiday effects, seasonality, etc.), and can accurately model count time series. We need to meet these modeling goals with an approach that can scale to large amounts of data for city-wide human mobility modeling.

Our approach meets all of these modeling requirements by preprocessing human mobility data into occupancy counts in space-time bins and then modeling and forecasting these counts with DGLMs. DGLMs additionally have very competitive forecast accuracy compared to competing modeling approaches [3]. Other related work does not meet all of these criteria; mechanistic and deep learning approaches often do not represent uncertainty and can be challenging to scale. Statistical models like Gaussian processes and extensions with complex dependencies can also be challenging to scale. Previous state space modeling approaches have in general been applied at a smaller scale than in this work, or to different spatiotemporal data that is not human mobility data. To the best of our knowledge, our application of DGLMs, and specifically the sparse count mixture variations [3, 30], to large scale mobility data modeling and forecasting is novel.

1.2 Contributions

In this work, we present a flexible and scalable approach to modeling human mobility data with dynamic generalized linear models. We use a preprocessing approach to convert raw spatiotemporal data to occupancy counts in space-time bins. Occupancy counts for each spatial bin are modeled with DGLMs and we present forecasting results for several spatial resolutions and forecast horizons; our forecast point accuracy does not degrade between fifteen-minute (1-step ahead) forecasts up to twenty-four hour (96-step ahead) forecast horizons. Across a variety of different levels and patterns of observed occupancy counts, we are able to accurately model and forecast with this common modeling approach. Finally, we present details on the scaling and implementation of our approach,

which is linear in the number of observations, spatial cells, and time bins. While our results focus on human mobility data, this work is applicable to other application areas for spatiotemporal data where accurate probabilistic modeling at scale is the main interest. Our approach supports fully online preprocessing, updating, and forecasting, and anomaly detection for aggregate behaviors.

2 Data and Application Goals

In this work, we focus on modeling trajectory data, where each observation consists of an agent ID, latitude, longitude, and timestamp. Observations can be unevenly sampled in time and space, as there is often sensor noise in the latitude/longitude values and IDs for agents may not be persistent. Additionally, human mobility data tends to be largely stationary, with movement occurring between stationary locations. Aggregate occupancy counts exhibit a large range of values, from very sparsely occupied cells to densely populated locations. Our modeling approach needs to address these challenges. Our goal for this work is to flexibly model large-scale human mobility data to support a variety of downstream applications, such as forecasting traffic flow and performing aggregate-level anomaly detection.

The results presented in this work use simulated trajectory data, following the generative process described in [27]. Specifically, the generative data model is parameterized by sequences of persistent locations that are associated with each agent in the synthetic data. The data generating process assumes that agents often: (1) have recurring behaviors in the same locations (e.g. home and work locations), (2) live and behave similarly to other agents, and (3) travel along roads between their persistent locations [27].

Persistent locations are randomly assigned to each agent using foundation data from [6, 19, 26], and the assignment of persistent locations encourages similarity among groups of agents (e.g. agents share the same office location or have nearby houses). Agents travel on roads between persistent locations; the synthetic data uses an Open Street Maps (OSM) road network from the OSMnx Python package [4], which results in a high frequency human mobility dataset. For our experiments below, we start with this simulated trajectory data sampled every 60 seconds for 10909 agents over a 28-day period (Figure 1). We do not assume specific knowledge of this data generating procedure in our results below.

3 Spatiotemporal Modeling with DGLMs

3.1 Data Preprocessing

To meet our goal of flexibly modeling human mobility data at scale, we first preprocess the raw trajectory data (ID, latitude, longitude, timestamp observations) into occupancy counts of unique agents in space-time bins. This discretization of space and time and reduction from individual agent-level trajectories to aggregate occupancy counts significantly reduces the volume of data and enables parallelization across relevant dimensions in the modeling. Additionally, converting the data to occupancy counts in discretized space-time bins provides robustness to sensor noise, uneven sampling rates of the data over time, and ID dropout or confusion. While discretizing the data over space and time and then aggregating into occupancy counts loses information at the individual observation level, this approach supports our goal of flexibly modeling at scale.

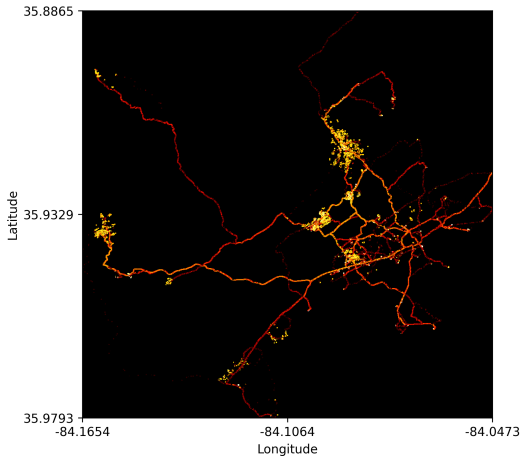


Figure 1: Synthetic human mobility data for 10909 agents over a 28-day time period. Brighter colors indicate a higher density of observations. The area of interest is large, resulting in both dense and sparse areas to model.

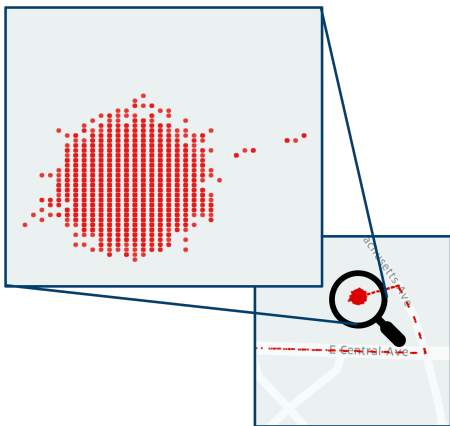


Figure 2: Simulated sensor noise from our synthetic human mobility data. We encode our mobility data into space-time bins to be robust to sensor noise.

Our preprocessing approach is outlined in Figure 3. First, we encode all latitude, longitude values into a spatial bin. In this work, we use S2 cells [9] to discretize space, though our approach is flexible to any choice of geohashing algorithm, including geohash [15] or H3 cells [8]. Next, we convert the raw timestamp values into temporal bins based on a chosen temporal resolution. This encoding of space and time can occur in parallel across observations and does not depend on the choice of spatial or temporal resolution (unless the geohashing algorithm takes longer to encode for finer resolutions). Finally, to form the occupancy counts, we count the number of unique agents in each space-time bin. We allow each agent to be in multiple spatial bins within the same time bin, though this assumption is not required; the first, last, mean, etc. location for each time bin could be selected instead. This procedure converts the raw trajectory observations into aggregate occupancy counts

over spatial and temporal bins for the entire dataset of interest in a scalable manner.

In this work, we consider three different spatial resolutions to demonstrate the flexibility of our approach; summary statistics for these spatial resolutions are given in Table 1. The temporal resolution for our results is 15-minute time bins, for a total of 2688 time bins over 28 days of synthetic data. Finer spatial and/or temporal resolutions decrease the loss of spatial information from discretizing the data, but increase the sparsity in the resulting occupancy counts and increase the sensitivity to noise. The degree of sparsity for each spatial cell increases at finer spatial resolutions (Figure 4), which also impacts the modeling. In Table 1, at finer spatial resolutions, more spatial cells require mixture models (DCMMs or DLMMs, described below) to model the observed occupancy counts.

3.2 Dynamic Generalized Linear Models

After converting the spatiotemporal data to occupancy counts in space-time bins, we model the resulting data with dynamic generalized linear models (DGLMs) [28, 29]. DGLMs are Bayesian state space models that can easily incorporate seasonality, covariates, and holiday effects. DGLMs have straightforward probabilistic forecast distributions, and are flexible to a variety of occupancy count distributions, from very high-levels of counts in each spatial cell, to very sparse cells. In this work, we fit one DGLM to each spatial cell in our area of interest.

Let y_t be the observed occupancy count at time bin t for one spatial cell. We assume that y_t comes from an exponential family distribution and define a DGLM to model this sequence of occupancy counts as follows [28]:

$$\begin{aligned} y_t &\sim \text{ExpFamily}(\eta_t) \\ \lambda_t &= \mathbf{F}_t' \boldsymbol{\theta}_t = g_{\text{link}}(\eta_t) \\ \boldsymbol{\theta}_t &= \mathbf{G}_t \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t, \\ \boldsymbol{\omega}_t &\sim (\mathbf{0}, \mathbf{W}_t), \end{aligned} \quad (1)$$

where $g_{\text{link}}(\cdot)$ is the link function for the chosen exponential family form, η_t is the natural parameter, $\boldsymbol{\theta}_t$ is the state vector, \mathbf{F}_t is a vector of known dynamic covariates, \mathbf{G}_t is a known state matrix, $\boldsymbol{\omega}_t$ is the evolution noise, and $\mathbb{E}(\boldsymbol{\omega}_t) = \mathbf{0}$, $\mathbb{V}(\boldsymbol{\omega}_t) = \mathbf{W}_t$. Dynamic covariates, such as point of interest information, holiday effects, etc. are incorporated into the \mathbf{F}_t vector, and seasonality is also incorporated via \mathbf{F}_t and \mathbf{G}_t . See Section A.2 and [28] for more details about model inference.

The \mathbf{W}_t matrices represent the evolution variance and are specified in terms of component discounting [28]. \mathbf{W}_t is partitioned into components for the trend, seasonal terms, and covariates, and each component includes a specific discount factor. Discount factors specify the level of change in the stochastic state evolution and how quickly the DGLM state vectors adapt to new information; discount factors close to 1 indicate very little stochastic noise in the state evolution, while discount factors closer to 0 indicate large variation in the respective state vector, allowing the DGLMs to adapt very rapidly to incoming information and “discount” the value of older observations [28]. In practice, discount factors are rarely set lower than about 0.9. Discount factors are the main hyper-parameters to

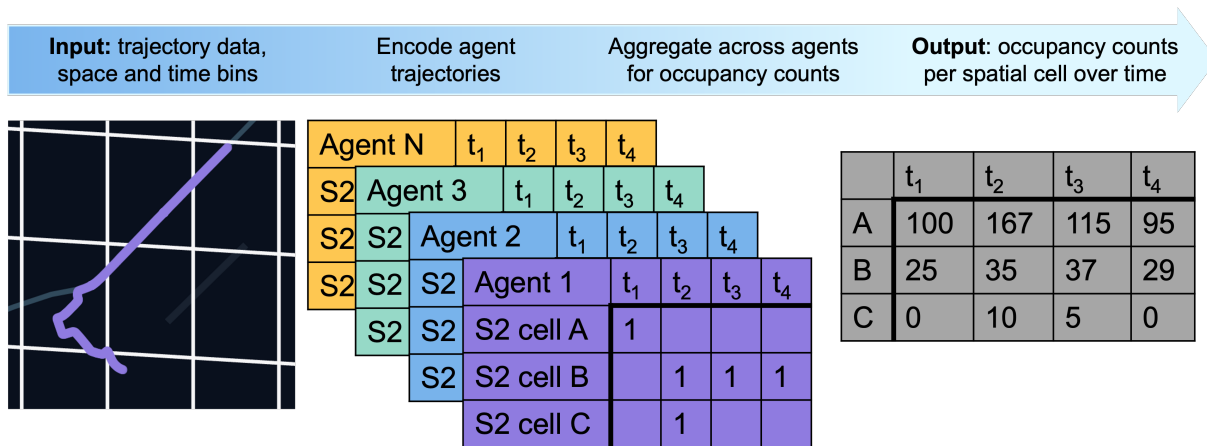


Figure 3: Preprocessing procedure to convert raw mobility data into occupancy counts in space-time bins. The encoding occurs in parallel across all observations, making this approach highly scalable.

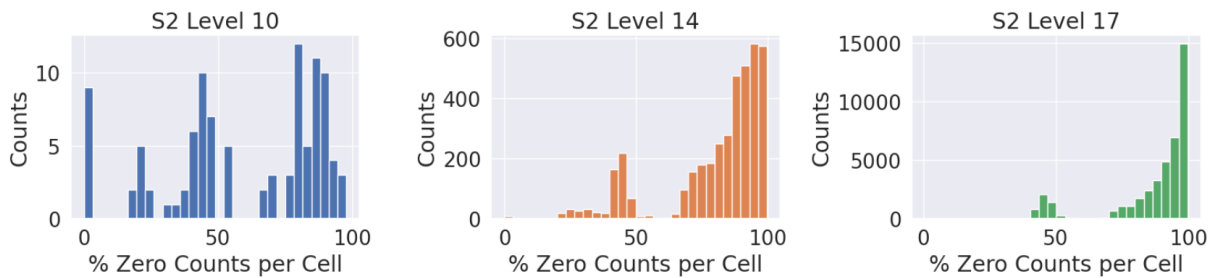


Figure 4: Sparsity (or percent of zero occupancy counts per spatial cell) for three different spatial resolutions. At fine spatial resolutions (e.g. S2 cell level 17), the majority of spatial cells have zero occupancy for more than 50% of the 15-min time bins over the month of our synthetic data.

tune for DGLMs and can be manually selected based on domain knowledge or tuned with a grid search.

In this work, we automatically choose the specific DGLM form based on the observed level of occupancy counts in the initial training data. Due to the choice of conjugate priors for the parameters in the DGLM, these models have closed-form forecast distributions, which makes forecasting efficient [28]. Spatial cells with high levels of occupancy counts (defined as a mean occupancy count over the training period of > 50) are fit with a normal dynamic linear model (DLM), where the occupancy counts are assumed to follow a normal distribution and the forecasts are t-distributed. For spatial cells with a mean occupancy count over the training data of < 50 , we fit a Poisson DGLM; forecasts are Negative Binomial distributed. We also use mixture model extensions to DGLMs to model and forecast the sparse occupancy counts that we observe in a variety of spatial cells. Dynamic count mixture models (DCMMs) [3] are a mixture of a Bernoulli DGLM (to model and predict zero counts) and a Poisson DGLM (to model and predict the non-zero counts) and can model a variety of different observed sparsity levels. DCMMs also include a random effects extension for over-dispersed count data. The forecast distribution for the DCMM is the product of Beta-Bernoulli distributed forecasts from the Bernoulli DGLM

and Negative Binomial forecasts from the Poisson DGLM. We fit DCMMs to spatial cells that have a mean < 50 and an observed percent of zero occupancy counts that is greater than 15%; as seen in Table 1, we fit DCMMs to many spatial cells across all spatial resolutions, but especially for finer resolutions. Additionally, for spatial cells that exhibit high average occupancy count levels, but still have a degree of sparsity $> 15\%$, we use dynamic linear mixture models (DLMMs) [30], which are a mixture of a Bernoulli DGLM and a normal DLM; forecast distributions are a mixture of samples from the Beta-Bernoulli and the t-distribution.

We include a trend term (local level or dynamic intercept) and daily seasonality in the DGLMs for all spatial cells. We also need to specify trend and seasonality discount factors for all DGLMs, random effect discount factors for the DCMMs and Poisson DGLMs, and a discount factor on the stochastic variance of the DLM and DLMM observations. Based on offline tuning experiments, we set the seasonality discount factor to 0.994 and the random effect discount factor or the stochastic variance discount factor to 0.9 for all spatial cells. We tune the trend discount factor using a grid search over the possible values of $[0.96, 0.97, 1.0]$ for each spatial cell individually and select the value that minimizes the mean absolute error of the forecast median over the first 3 days of data (results below are

presented for the following 25 days of simulated data). Additional tuning of all discount factors via a grid search is straightforward, though does add to the computation time.

DGLM updates are linear in the number of timesteps and do not require expensive Markov-chain Monte Carlo or sequential Monte Carlo, which makes these models efficient for inference and forecasting at scale in our application. Additionally, this updating occurs in an online, sequential manner, enabling our approach to support streaming data settings. Updating and forecasting only require the parameter values from the previous timestep, making the DGLMs efficient in terms of memory, as well. We fit one model per spatial cell and these models are updated in parallel across space.

3.3 Forecast Metrics

We evaluate how well our modeling approach fits the observed occupancy counts in terms of both point and uncertainty forecast metrics. We evaluate the accuracy of our point forecasts using root-mean squared error (RMSE), mean absolute error (MAE), and zero-adjusted percent error (ZAPE); ZAPE explicitly evaluates how well our models predict zero vs. non-zero occupancy counts. The specific form of ZAPE used here follows [30] and the loss function is defined for a non-negative count y and forecast f as:

$$\mathcal{L}_{ZAPE} = \mathbf{1}(y = 0) \frac{f}{(1+f)} + \mathbf{1}(y > 0) \frac{|y-f|}{y}, \quad (2)$$

where $\mathbf{1}(\cdot)$ is the indicator function. The RMSE and MAE metrics are on the same scale as the data (occupancy counts) and ZAPE is a percent error metric. Point forecast metrics below are calculated using the forecast median. We also use coverage to evaluate the calibration of our forecast distributions; for well-calibrated uncertainty estimates, we expect our empirical coverage to be close to the theoretical coverage value.

3.4 Scaling Summary

Our preprocessing approach to convert the raw trajectory data observations to occupancy counts is linear in the number of observations (or linear in the number of agents if each agent has approximately the same number of observations). Modeling and forecasting with DGLMs is linear in the number of time bins, with 1 DGLM fit to each spatial cell. In summary, our overall approach is linear in the number of observations, time bins, and spatial bins.

3.5 Extensions

There are several straightforward extensions to our approach that would only require updates to the preprocessing procedure, but not to the overall modeling approach. While we focus on modeling occupancy counts in this work, extending to the common setting of modeling transitions between spatial cells, or origin-destination matrices [2], is trivial. Instead of counting the number of agents in each space-time bin in our preprocessing procedure, we would calculate transitions, or flow counts, of the number of agents that transition from spatial cell i to spatial cell j at time t and would fit a DGLM for every observed transition between spatial cells (including self-transitions). This would greatly increase the number of DGLMs to fit, which is why we did not consider this setting here,

but otherwise would require no updates to our modeling approach. We would expect similar modeling results to those shown below, with the observations of occupancy counts simply being updated to flow counts. In addition to increasing the number of DGLMs, modeling transitions instead of occupancy counts would increase the level of sparsity in each observed transition; however, DCMMs could still be applied to capture varying levels of sparsity in these transition counts.

An advantage of DGLMs for modeling human mobility data is the ease of incorporating external, dynamic, information into the DGLMs directly. Here, we only use a trend term and daily seasonality for our synthetic data results, though we could easily incorporate holiday effects, point of interest information, etc. into the models. We also only tune the DGLM discount factor hyper-parameters over a limited grid to speed-up computation, but could increase the degree of hyper-parameter tuning to further improve upon forecast accuracy results presented below.

Finally, our approach requires a choice of temporal and spatial resolutions. There is a trade-off between finer resolutions and the degree of sparsity for each spatial cell that is modeled (Figure 4); while DCMMs can accurately model and forecast count time series with a variety of levels of sparsity, the more sparse the time series becomes, the more challenging it is to accurately predict where the non-zero observations will occur [30]. To address this challenge, we could model over multiple spatial resolutions, so that the spatial cells were not uniform in area but were variably sized based on a base level of occupancy counts; this would only require changes to our preprocessing, but not our modeling approach. Alternatively, we could form multi-scale DGLMs over multiple spatial and/or temporal resolutions [3, 11, 30] to share information at more aggregate resolutions with the finer-grained resolutions; this information could be added as an additional covariate in our DGLMs with no further updates to the modeling approach. Similarly, we could incorporate spatial dependence between cells by using the forecasted occupancy count of one cell as a covariate in the DGLM for another spatial cell.

3.6 Anomaly Detection

The focus of this work is on accurate probabilistic modeling and forecasting of occupancy counts, however our approach supports several additional downstream applications, including anomaly detection. The specific manner in which to use the DGLMs for anomaly detection depends on the types of anomalies that need to be detected; modeling occupancy counts with DGLMs as presented here supports anomaly detection over aggregate changes in agent behavior, which is reflected in higher or lower occupancy counts than expected in a spatial cell. For example, in a transportation traffic monitoring application, higher occupancy counts than expected could represent a traffic jam, while lower counts than expected could represent a road closure. Anomalous occupancy counts can be detected by comparing the quantile or likelihood of the observed occupancy count at time t under the one-step ahead forecast distribution to the “expected” count at time t . This “expected” value could be a multi-step ahead forecast from further back in time or from another DGLM that does not adapt as quickly to current data.

Table 1: Spatial resolution and modeling details for different S2 cell levels; S2 cell information from [9]. Finer spatial resolutions exhibit more sparsity, requiring the majority of spatial cells to be DCMMs or DLMMs, as compared with the coarser cells.

S2 Cell Resolution	Avg. Cell Area	Avg. Cell Edge Length	# Spatial Cells / # Models	# DLMs	# DLMMs	# Poisson DGLMs	# DCMMs
10	81.07 km ²	8.8 km	103	9	30	0	64
14	0.32 km ²	550 m	3938	7	41	1	3889
17	4948 m ²	70 m	42038	0	35	1	42002

DGLMs are probabilistic models, which allows for uncertainty estimates to propagate into the anomaly detection, too.

4 Results

4.1 Forecast Metrics and Examples

Across a variety of spatial resolutions and forecast metrics, DGLMs accurately model the observed occupancy counts and the uncertainty estimates of the forecast distributions are well-calibrated (Table 2). The RMSE and MAE forecast metrics are on the same scale as the data; on average, the 1-step ahead forecasts across spatial cells are only a handful of occupancy counts off from the observations, and the average percent error is also low for the finer resolution spatial cells. The coverage of the DGLMs at 90% tends to be > 90%, indicating that the uncertainty estimates are slightly wider than they need to be; in practice, over-coverage is often preferred to under-coverage. Additional tuning of the discount factors, specifically the discount factor for the random effects in the Poisson DGLM and DCMM [3], could improve the calibration of the uncertainty estimates. However, with a single modeling framework, we are able to model a variety of different occupancy counts, demonstrating the flexibility of our approach.

The magnitude of the observed occupancy counts does impact the forecast accuracy of the models (Figure 5); the DLMs are used to model spatial cells with higher occupancy counts and the forecast medians have more potential to differ from the observed occupancy counts. On the other hand, the DCMMs are used for lower count spatial cells, leading to lower forecast metric values. The finer the spatial resolution, the lower the occupancy counts for each spatial cell overall, which is why the point forecast metrics in Table 2 improve with finer spatial resolutions. However, at finer spatial resolutions, sparsity also increases. When the degree of sparsity for each spatial cell is particularly high, e.g. > 70% of counts for a single spatial cell are zero, it is challenging to predict when the non-zero occupancy counts will occur. Modeling could be improved in these settings with additional covariates; for example, point-of-interest information could aid in predicting a zero vs. non-zero count for these sparse cells. The level of noise in the data is also a consideration; as the spatial resolution becomes finer, it is more likely that sensor noise will cause occupancy counts to be split across spatial cell boundaries. This can be addressed in several ways; for example, smoothing the raw trajectory data, only encoding one observation per agent per time bin, or considering multi-scale approaches over multiple spatial resolutions [11]. In practice, the specific application goals will guide the choice of spatial and temporal resolutions for modeling the occupancy counts.

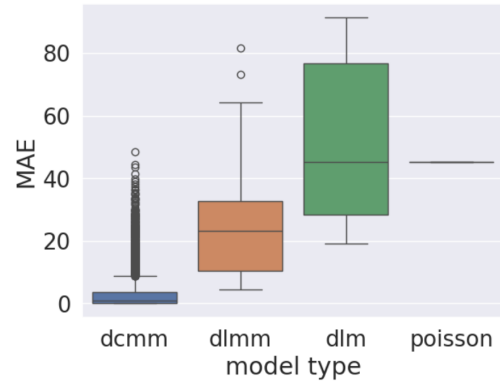


Figure 5: Mean absolute error (across spatial cells) by model type for S2 cell level 14. DLMs are used to model higher occupancy count cells and can have larger forecast errors than the DCMMs, used for low and sparse occupancy count cells.

DGLMs accurately model and forecast a variety of levels of occupancy counts, including high counts fit with a DLM, moderate counts with sparsity fit with a DLMM and sparse, lower levels of counts fit with a DCMM (Figure 6). Across these example spatial cells, the forecast means closely track the observed occupancy counts and the 90% credible intervals are well calibrated, as seen across all spatial cells in Table 2. The flexibility of the DCMMs and DLMMs, in particular, to model spatial cells with a high-level of sparsity (or 0 occupancy counts), is critical, as across spatial resolutions, we observe many spatial cells with some degree of sparsity, varying from minimal sparsity to more than 80% sparsity (Figure 4). In our synthetic data, this sparsity occurs as a result of zero counts in a given spatial cell, but the DGLMs also easily handle missing data by not updating for timesteps that have no observations. The majority of spatial cells in our synthetic data exhibit strong daily seasonality, which is explicitly incorporated into each of the DGLMs and could be expanded to include seasonality at additional resolutions, like weekly or monthly.

Spatial cells with high levels of counts and high levels of sparsity are in general challenging to model; these cells contribute to the larger forecast errors for S2 level 10 spatial cells in Table 2. An example spatial cell is shown in Figure 7. This spatial cell exhibits relatively high, non-zero counts (82.2 on average) and a large degree of sparsity; 19% of time steps have zero occupancy counts. A DLMM is automatically selected to model this spatial cell, and the

Table 2: Forecast accuracy metrics across spatial cells for 1-step ahead forecasts. Mean and standard deviation across all cells is shown for the forecast median. RMSE: root-mean squared error, MAE: mean absolute error, ZAPE: zero-adjusted percent error. Lower is better for RMSE, MAE, and ZAPE and closer to the nominal coverage (e.g. 80%, 90% or 95%) indicates better calibrated coverage.

S2 Cell Resolution	RMSE	MAE	ZAPE	80% Coverage	90% Coverage	95% Coverage	# of Models
10	65.62 ± 87.24	36.17 ± 57.73	63.53 ± 137.85	92.40 ± 4.48	94.87 ± 3.35	96.04 ± 2.98	103
14	9.18 ± 14.68	3.81 ± 7.15	22.87 ± 40.51	94.37 ± 5.99	95.87 ± 4.62	96.62 ± 4.16	3938
17	1.10 ± 2.15	0.37 ± 1.03	8.49 ± 8.99	97.32 ± 2.46	98.38 ± 1.72	98.97 ± 1.35	42038

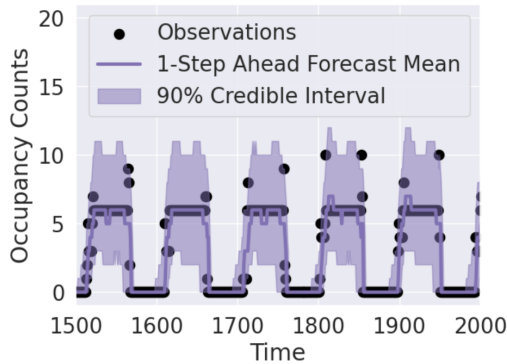


Figure 6: One-step ahead forecast mean and credible intervals for a DCMM fit to a spatial cell with low, sparse occupancy counts at S2 cell level 14.

model has relatively high forecast error. This cell is challenging to model because of the sparsity and because the non-zero counts are not normally distributed. Alternatively, we could use a Poisson DGLM to model this spatial cell (Figure 8); while this improves the MAE and ZAPE metrics, coverage is significantly worse. Even allowing for over-dispersion, the Poisson distribution has only one parameter for the mean and variance. In these high count settings, the Poisson DGLM significantly under-estimates the uncertainty in the occupancy counts. Alternatively, we can tune the trend and stochastic variance discount factors over a larger grid of values to improve the forecast accuracy (Figure 9). The optimal discount factors are much lower than the values selected in Figure 7; 0.9 for the discount factor for the trend and 0.6 for the stochastic variance in Figure 9 compared to 0.96 for the discount factor for the trend and 0.9 for the stochastic variance in Figure 7. Lower values for the stochastic variance increase how much over-dispersion is expected in the counts, leading to the improved uncertainty estimates. Additionally, the lower discount factor for the trend allows the model to adapt more quickly to the incoming data, leading to a slight improvement in the forecast accuracy. Finally, the forecast accuracy for these cells could likely be further improved by additional preprocessing to make the normal distribution assumption for the non-zero counts more valid, or considering different distributions than a mixture of a Bernoulli DGLM and a normal DLM via the DLMM. However, our results here demonstrate the flexibility and feasibility of our approach to a variety of choices of spatial and temporal resolutions, and across different DGLM model types.

4.2 Multi-Step Ahead Forecasting

Multi-step ahead forecasting is also straightforward with DGLMs, and the DGLM point forecasts remain accurate across longer-range forecast horizons (Table 3). Our temporal resolution is 15 minutes for the mobility data, and we compare results for 1 timestep ahead to forecast results for 1 hour (4 steps ahead), 6 hours (24 steps ahead) and 24 hours (96 steps ahead). Forecast accuracy across metrics is comparable at all of these forecast horizons, demonstrating the ability of DGLMs to be used for long-range forecasting for human mobility. Results below are for marginal, rather than path, forecasts [3, 28]. The coverage is lower at multi-step ahead forecast horizons because we do not propagate the state vector uncertainty (via the discount factors) at longer than 1-step ahead forecast horizons in these results, though that additional uncertainty could easily be incorporated to the multi-step ahead forecasts. We could alternatively use path forecasts to improve the uncertainty estimates, at the expense of longer computation times than the marginal forecasts shown here. The long-range forecast accuracy benefits from incorporating the strong daily seasonality explicitly into the DGLMs, which is very straightforward to do with these models. Other relevant external information could also be easily incorporated to further improve forecast accuracy as needed.

Figure 10 show twenty-four hour ahead forecast results (96-steps ahead) for a DLMM fit to a relatively sparsely occupied spatial cell at S2 level 14. The forecasts at this long forecast-horizon are accurate, do not degrade much from the 1-step ahead forecasts, and still have well calibrated forecast uncertainty. These results demonstrate the flexibility of the mixture models (DLMMs and DCMMs), in particular, to accurately model very sparse occupancy counts and the ability of DGLMs to accurately forecast over long time horizons; important qualities when using these models for a variety of downstream applications.

4.3 Robustness

The synthetic data used in the experiments above has very regular seasonality in the occupancy counts. However, our approach is flexible to both different types of seasonality and less regularity in the seasonality. To demonstrate this feature of the DGLMs, we also include robustness results that focus on adding noise to the period of the daily seasonality in the occupancy counts. For each spatial cell at S2 level 14, we randomly sample a slight offset independently for each day to shift the seasonal pattern slightly forward or slightly backward in time (we also include no shift as an option to randomly select). This additional noise in the data is meant to represent real world data, where there may be regular seasonality, but the period

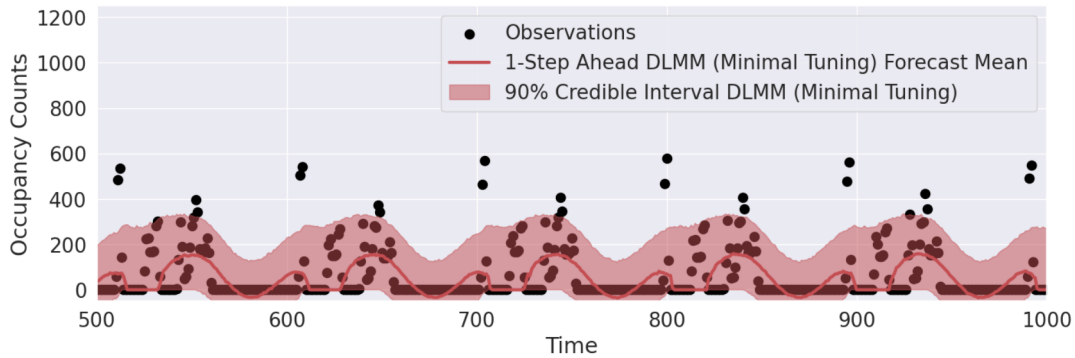


Figure 7: One-step ahead forecast mean and credible intervals for a S2 level 10 cell modeled with a DLMM. This cell is challenging to model, due to the high number of zero counts and the high values of non-zero occupancy counts. For the base DLMM, where the only discount factor tuned is the trend term over a grid search of [0.96, 0.97], the RMSE for the 1-step ahead forecast mean is 109.75, the MAE is 70.91, the ZAPE is 984.40, and the 90% empirical coverage is 95.31%.

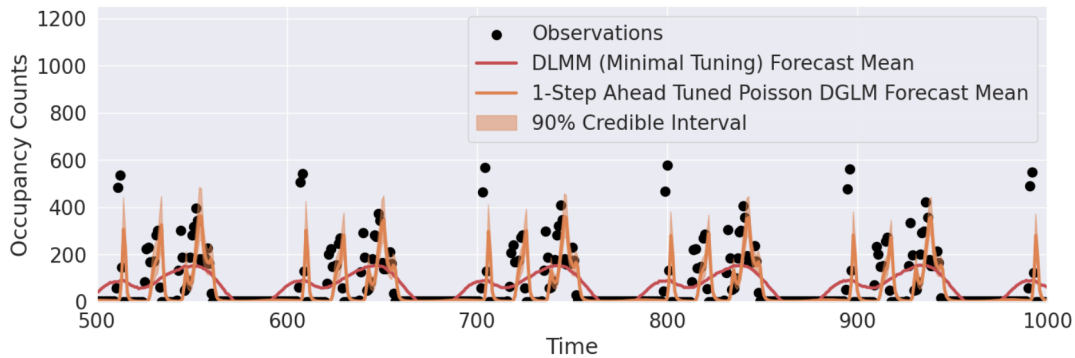


Figure 8: One-step ahead forecast mean and credible intervals for the same S2 level 10 cell in Figure 7, this time modeled with a Poisson DGLM. The trend discount factor is tuned over [0.9, 0.93, 0.95, 0.96, 0.97, 0.98, 0.99, 0.999, 1] and the random effect discount factor over [0.6, 0.7, 0.8, 0.85, 0.9, 0.95, 0.99, 1]. While the point forecast metrics improve slightly with this Poisson DGLM, the uncertainty estimates are much worse. The RMSE for the 1-step ahead forecast mean is 117.93, the MAE is 62.12, the ZAPE is 221.91, and the 90% empirical coverage is 50.0%.

Table 3: Forecast accuracy metrics across S2 level 14 spatial cells for multi-step ahead forecasts. Table format as in Table 2.

Forecast Steps Ahead	Time Ahead	RMSE	MAE	ZAPE	80% Coverage	90% Coverage	95% Coverage
1	15 min	9.18 ± 14.68	3.81 ± 7.15	22.87 ± 40.51	94.37 ± 5.99	95.87 ± 4.62	96.62 ± 4.16
4	1 hour	11.20 ± 18.45	4.76 ± 8.88	23.02 ± 38.74	73.22 ± 4.37	74.39 ± 3.58	75.04 ± 3.21
24	6 hours	9.06 ± 15.00	3.68 ± 7.35	20.79 ± 38.44	73.26 ± 3.51	74.17 ± 2.98	74.72 ± 2.67
96	24 hours	9.35 ± 14.87	3.91 ± 7.31	23.27 ± 38.71	69.46 ± 4.56	70.91 ± 3.40	71.49 ± 3.08

of the seasonality is not as regular as the synthetic data we work with here; each day the occupancy count peak might shift sooner or later in the day, and this pattern is random.

We add this random noise in all spatial cells at S2 level 14 for each day of data, over four possible levels of offset: up to 0.5 hours, up to 1 hour, up to 2 hours, and up to 3 hours sooner or later than expected. For all settings, no offset is also a possible noise setting. We then fit the DGLMs as described above to all spatial cells and compare to a baseline model that just uses the occupancy count $t -$

24 hours ago to predict the occupancy count at time t . The DGLMs still have daily seasonality in the model and we do not change any experimental settings except to retune the trend discount factor for the noisy data. Except for retuning and fitting, the modeling setup is identical to the results above without seasonal noise.

In addition to providing uncertainty estimates, the DGLMs are able to flexibly adapt to noise in the seasonal pattern, across different levels of noise offset, while the simple baseline models are not. Forecast metrics for these experiments are shown in Table 4

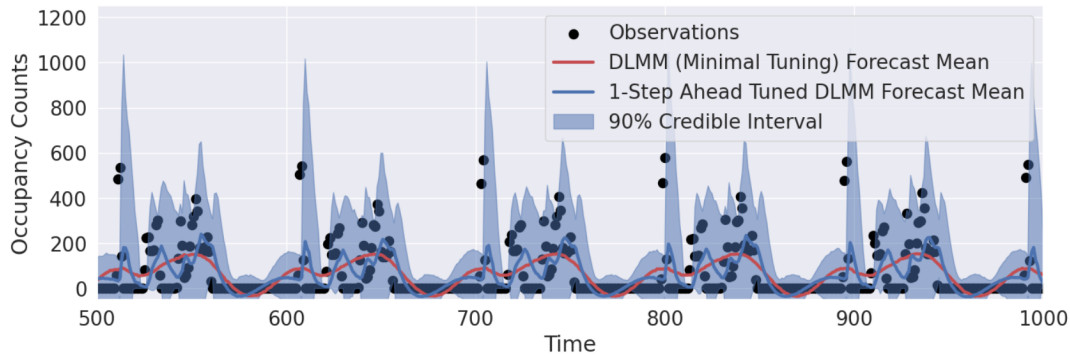


Figure 9: One-step ahead forecast mean and credible intervals for the same S2 level 10 cell in Figure 7 for a tuned DLMM. The trend discount factor is tuned over $[0.9, 0.93, 0.95, 0.96, 0.97, 0.98, 0.99, 0.999, 1]$ and the stochastic variance discount factor over $[0.6, 0.7, 0.8, 0.85, 0.9, 0.95, 0.99, 1]$. The uncertainty estimates are slightly better in this case, compared to the coverage in Figure 7 and the point forecast metrics are similar. The RMSE for the 1-step ahead forecast mean is 114.42, the MAE is 77.77, the ZAPE is 956.41, and the 90% empirical coverage is 93.0%

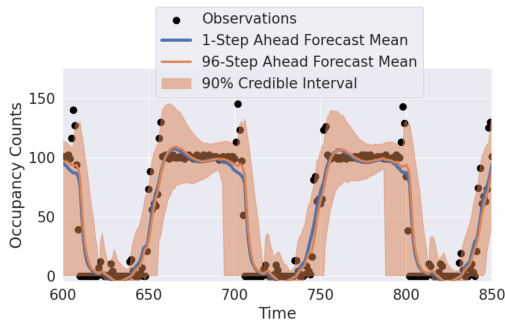


Figure 10: Twenty four-hour ahead (96-step ahead) forecast mean and credible intervals, compared to the 1-step ahead forecast mean, for a DLMM fit to a spatial cell at S2 level 14.

and example forecasts are shown in Figure 11 and Figure 12. DGLM forecast accuracy remains the same throughout these offsets, while the baseline forecast accuracy degrades with larger offsets. These results demonstrate the utility of our approach on realistic data, that is expected to have more variation in seasonality.

4.4 Scaling and Implementation

Results shown above were for $\sim 11k$ agents over one month of data. This simulated data was 9.77 GB, compressed and at rest and results were run on a machine with 48 CPU cores. Computation and memory requirements for the $\sim 11k$ agents are shown in Table 5; our preprocessing procedure significantly reduces the size of the input data to model. Additionally, preprocessing time is approximately constant across spatial resolutions, as expected, and modeling time increases linearly with the number of spatial cells to model; we fit and forecast DGLMs across $\sim 4k$ spatial cells in just over 1 hour.

We have also applied our approach at a much larger scale and are still able to preprocess and model the data efficiently. With our approach, the preprocessing time depends on the number of raw observations and not the choice of spatial or temporal resolutions

Table 4: Forecast accuracy metrics across S2 level 14 spatial cells with noise added to the daily seasonal period. For example, an offset of 1 hour means that the possible daily seasonal periods were 23, 23.5, 24, 24.5, and 25 hours; this period was randomly chosen for each day for each spatial cell. Table format as in Table 2.

Offset (hrs.)	Model	RMSE	MAE	ZAPE
0.5	DGLM	9.26 ± 14.91	3.86 ± 7.32	23.23 ± 42.21
	Baseline	10.07 ± 15.79	3.57 ± 6.00	26.04 ± 29.80
1	DGLM	9.38 ± 15.33	3.94 ± 7.62	23.39 ± 41.89
	Baseline	11.51 ± 18.92	4.45 ± 8.10	28.26 ± 34.39
2	DGLM	9.74 ± 16.84	4.08 ± 8.36	23.19 ± 47.92
	Baseline	12.65 ± 21.24	5.25 ± 9.89	29.35 ± 39.00
3	DGLM	9.93 ± 17.31	4.15 ± 8.89	22.48 ± 47.39
	Baseline	13.55 ± 22.99	5.81 ± 11.34	31.09 ± 54.95

for the modeling. Our preprocessing also significantly reduces the size of the data to model. We applied this approach to additional synthetic data with **484 million** observations which was 472 GB of data compressed and at rest. After transforming this data to occupancy counts with a spatial resolution of S2 cell level 17 and temporal resolution of 1 hour, the data size was reduced to **0.582 GB** compressed. Preprocessing the 484 million observations took 10 hours on a 384-core cluster. We then fit 397352 DGLMs on the same size cluster in 10 hours; each DGLM updated and forecasted in an average of 2.5 seconds for 672 timesteps, or approximately 4 ms per DGLM per timestep to update and forecast.

Our approach was implemented in PySpark [12], with Apache Sedona [13, 31] to encode the data into S2 cells and PyBats [16] to fit the DGLMs. Further efficiencies are possible with additional tuning of the Spark cluster and improvements to the implementation of the DGLMs. However, our results show the feasibility of our approach

Table 5: Computation time and memory requirements for the results shown on ~11k agents for 28 days; the input data was 9.77 GB compressed and at rest. Preprocessing significantly reduces the size of the data to model and preprocessing time is the same across spatial resolutions; modeling time scales linearly with the number of spatial cells/models.

S2 Cell Resolution	Preprocessing Time (Min.)	Modeling Time (Min.)	Total Time (Min.)	Preprocessing Output (MB)	# of Models
10	15.17	18.84	34.01	0.32	103
14	16.45	66.48	82.94	7.0	3938
17	16.49	526.92	543.41	49.0	42038

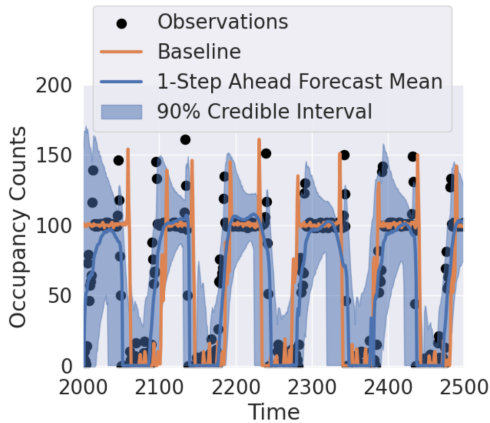


Figure 11: One-step ahead forecast mean and credible intervals, compared to the seasonal baseline, for a DLMM fit to a spatial cell at S2 level 14 with seasonal offset up to 3 hours. The DLMM is better able to adapt to slightly varying seasonal periods.

for very large scale datasets both in terms of the preprocessing and the modeling; our preprocessing scales linearly with the number of observations and the modeling scales linearly with the number of space-time bins. Our approach is very well-suited to a distributed computing environment, as each model only operates on a small portion of the data and is independent of all other models.

5 Summary and Future Work

In this work, we present a scalable and flexible approach to modeling and forecasting human mobility data. Our approach encodes the spatiotemporal data as occupancy counts in space-time bins and models these counts with DGLMs. DGLMs provide efficient inference, probabilistic modeling and forecasting, flexibility to a range of observed occupancy count levels (including very sparse occupancy counts and over-dispersed counts), and can easily incorporate external information, such as seasonality. We present forecast results on synthetic human mobility data for ~11k agents; DGLMs maintain high forecast accuracy for a range of spatial resolutions and for long forecast horizons. Finally, we show scaling results for much larger synthetic data consisting of hundreds of millions of observations to demonstrate the scalability of our approach.

Our approach is flexible enough to support a variety of downstream applications that require accurate, probabilistic forecasting

and modeling. For example, the forecast results could support traffic or transportation management, location demand forecasting, and urban planning. Additionally, the flexibility of our approach makes it applicable to other types of spatiotemporal data beyond human mobility, for example, our approach could aid in modeling animal movements at scale.

There are several directions for future work and extensions of our approach. One main direction, which would be largely application driven, is to incorporate additional information into the DGLMs. For example, the mobility data may include holiday effects that need to be accounted for, different patterns of seasonality, or we may wish to incorporate dynamic covariates for each spatial cell, related to point of interest information, for example. We could easily extend our approach to modeling flows of agents between spatial cells to better support movement or transportation related applications.

On the modeling side, a main direction for future work involves improving the modeling of very sparse spatial cells, especially cells that are both sparse and have large non-zero counts, as in Figure 7. Although DGLMs and DLMMs can accurately model a variety of sparse count levels, extremely sparse counts are challenging to model and forecast well; sparse time series modeling is a challenge in many different settings [30]. One solution would be to model at multiple resolutions or uneven spatial resolutions, to prevent any spatial cells from being too sparse. Another solution would be to share information over multiple spatial resolutions (and/or temporal resolutions) with a multi-scale approach [3, 11, 30].

Another area for future work is faster tuning of the discount factors than a grid search-based approach. For the majority of spatial cells, more fine-grained tuning of all discount factors improves forecast accuracy. However, with a grid search, we have to fit a DGLM for each combination of discount values to tune; it can be expensive to tune over too many unique values. Using a coarse approximation to this optimization problem could improve the forecast accuracy while also improving the scalability of our approach.

Acknowledgments

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number 140D0423C0046. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

- [1] Federico Amato, Fabian Guignard, Sylvain Robert, and Mikhail Kanevski. 2020. A Novel Framework for Spatio-temporal Prediction of Environmental Data using Deep Learning. *Scientific Reports* 10, 1 (2020), 22243. <https://doi.org/10.1038/s41598-020-79148-7>
- [2] Hugo Barbosa, Marc Barthelemy, Gourab Ghoshal, Charlotte R. James, Maxime Lenormand, Thomas Louail, Ronaldo Menezes, José J. Ramasco, Filippo Simini, and Marcello Tomasini. 2018. Human mobility: Models and Applications. *Physics Reports* 734 (2018), 1–74. <https://doi.org/10.1016/j.physrep.2018.01.001> Human mobility: Models and applications.
- [3] Lindsay R. Berry and Mike West. 2020. Bayesian Forecasting of Many Count-Valued Time Series. *Journal of Business and Economic Statistics* 38 (2020), 872–887. <https://doi.org/10.1080/07350015.2019.1604372>
- [4] Geoff Boeing. 2017. OSMnx: New Methods for Acquiring, Constructing, Analyzing, and Visualizing Complex Street Networks. *Computers, Environment and Urban Systems* 65 (2017), 126–139.
- [5] Xi Chen, David Banks, and Mike West. 2019. Bayesian Dynamic Modeling and Monitoring of Network Flows. *Network Science* 7 (2019), 292–318. <https://doi.org/10.1017/nws.2019.10>
- [6] OSM Contributors. 2017. Planet dump retrieved from <https://planet.osm.org>. <https://www.openstreetmap.org>.
- [7] Cecilia Cordeiro da Silva, Clarisse Lins de Lima, Ana Clara Gomes da Silva, Eduardo Luiz Silva, Gabriel Souza Marques, Lucas Job Brito de Araújo, Luiz Antônio Albuquerque Júnior, Samuel Barbosa Jatobá de Souza, Maira Araújo de Santana, Juliana Carneiro Gomes, Valter Augusto de Freitas Barbosa, Anwar Musah, Patty Kostkova, Wellington Pinheiro dos Santos, and Abel Guilhermino da Silva Filho. 2021. Covid-19 Dynamic Monitoring and Real-Time Spatio-Temporal Forecasting. *Frontiers in Public Health* 9 (2021). <https://doi.org/10.3389/fpubh.2021.641253>
- [8] H3 Developers. 2024. H3. <https://h3geo.org/>.
- [9] S2 Developers. 2024. S2 Geometry Library. <http://s2geometry.io/>.
- [10] Azadeh Emami, Majid Sarvi, and Saeed Asadi Bagloee. 2020. Short-term Traffic Flow Prediction Based on Faded Memory Kalman Filter Fusing Data from Connected Vehicles and Bluetooth Sensors. *Simulation Modelling Practice and Theory* 102 (2020), 102025. <https://doi.org/10.1016/j.simpat.2019.102025> Special Issue on IoT, Cloud, Big Data and AI in Interdisciplinary Domains.
- [11] Thais C. O. Fonseca and Marco A. R. Ferreira. 2017. Dynamic Multi-scale Spatiotemporal Models for Poisson Data. *J. Amer. Statist. Assoc.* 112, 517 (2017), 215–234. <https://doi.org/10.1080/01621459.2015.1129968> arXiv:<https://doi.org/10.1080/01621459.2015.1129968>
- [12] Apache Software Foundation. 2024. Apache PySpark. <https://spark.apache.org/docs/latest/api/python/index.html>.
- [13] Apache Software Foundation. 2024. Apache Sedona. <https://sedona.apache.org/1.6.0/>.
- [14] Trevor J. Hefley, Mevin B. Hooten, Ephraim M. Hanks, Robin E. Russell, and Daniel P. Walsh. 2017. Dynamic Spatio-temporal Models for Spatial Data. *Spatial Statistics* 20 (2017), 206–220. <https://doi.org/10.1016/j.spasta.2017.02.005>
- [15] Hiroaki Kawai. 2014. Python Geohash. <https://pypi.org/project/python-geohash/>.
- [16] Isaac Lavine and Andrew Cron. 2021. PyBats: A Python Package for Bayesian Analysis of Time Series and Bayesian Forecasting. <https://pypi.org/project/pybats/>.
- [17] Yanchi Liu, Chuanren Liu, Xinjiang Lu, Mingfei Teng, Hengshu Zhu, and Hui Xiong. 2017. Point-of-Interest Demand Modeling with Human Mobility Patterns. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Halifax, NS, Canada) (KDD '17)*. Association for Computing Machinery, New York, NY, USA, 947–955. <https://doi.org/10.1145/3097983.3098168>
- [18] Luis Leon Ojeda, Alain Y. Kibangou, and Carlos Canudas de Wit. 2013. Adaptive Kalman Filtering for Multi-Step Ahead Traffic Flow Prediction. In *2013 American Control Conference*. 4724–4729. <https://doi.org/10.1109/ACC.2013.6580568>
- [19] Oak Ridge National Laboratory (ORNL) and Federal Emergency Management Agency (FEMA) Geospatial Response Office. 2024. USA Structures. <https://gis-fema.hub.arcgis.com/pages/usa-structures>.
- [20] Henry R. Scharf, Mevin B. Hooten, Devin S. Johnson, and John W. Durban. 2018. Process Convolution Approaches for Modeling Interacting Trajectories. *Environmetrics* 29, 3 (2018), e2487.
- [21] Gürkan Solmaz and Damla Turgut. 2019. A Survey of Human Mobility Models. *IEEE Access* 7 (2019), 125711–125731. <https://doi.org/10.1109/ACCESS.2019.2939203>
- [22] Xuan Song, Hiroshi Kanasugi, and Ryosuke Shibasaki. 2016. DeepTransport: Prediction and Simulation of Human Mobility and Transportation Mode at a Citywide Level. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (New York, New York, USA) (IJCAI'16)*. AAAI Press, 2618–2624.
- [23] Aike Steentoft, Bu-Sung Lee, and Markus Schlöpfer. 2023. Quantifying the Uncertainty of Mobility Flow Predictions Using Gaussian Processes. *Transportation* (2023). <https://doi.org/10.1007/s11116-023-10406-z>
- [24] Jonathan R. Stroud, Peter Müller, and Bruno Sansó. 2001. Dynamic Models for Spatiotemporal Data. *Journal of the Royal Statistical Society (Ser. B)* 63, 4 (2001), 673–689.
- [25] Claudia Tebaldi, Mike West, and Alan F. Karr. 2002. Statistical Analyses of Freeway Traffic Flows. *Journal of Forecasting* 21 (2002), 39–68. <http://www.stat.duke.edu/mw/MWextrapubs/Tebaldi2002.pdf>
- [26] Gautam S. Thakur, Budhendra L. Bhaduri, Jesse O. Piburn, Kelly M. Sims, Robert N. Stewart, and Marie L. Urban. 2015. PlanetSense: a real-time streaming and spatio-temporal analytics platform for gathering geo-spatial intelligence from open source data. In *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems (Seattle, Washington) (SIGSPATIAL '15)*. Association for Computing Machinery, New York, NY, USA, Article 11, 4 pages. <https://doi.org/10.1145/2820783.2820882>
- [27] Athanasios Tsiligkaridis, Zhongheng Li, and Elizabeth Hou. 2024. Encoding Agent Trajectories as Representations with Sequence Transformers. *ICLR 2024 Machine Learning for Remote Sensing Workshop* (2024). https://ml-for-rs.github.io/iclr2024/camera_ready/papers/62.pdf
- [28] Mike West and Jeff Harrison. 1997. *Bayesian Forecasting and Dynamic Models* (2nd ed.). Springer. <http://www.stat.duke.edu/mw/West&HarrisonBook/>
- [29] M. West, P. Jeff Harrison, and Helio S. Migon. 1985. Dynamic Generalised Linear Models and Bayesian Forecasting (with Discussion). *J. Amer. Statist. Assoc.* 80 (1985), 73–97. <http://www.stat.duke.edu/mw/MWextrapubs/West1985a.pdf>
- [30] Anna K. Yanchenko, Di Daniel Deng, Jinglan Li, Andrew J. Cron, and Mike West. 2023. Hierarchical Dynamic Modelling for Individualized Bayesian Forecasting. *Journal of the Royal Statistical Society (Ser. C)* 72 (2023), 144–164. <https://doi.org/10.1093/jrsssc/qlad002> arXiv:2101.03408.
- [31] Jia Yu, Zongsi Zhang, and Mohamed Sarwat. 2019. Spatial data management in apache spark: the GeoSpark perspective and beyond. *Geoinformatica* 23, 1 (2019), 37–78. <https://doi.org/10.1007/s10707-018-0330-9>
- [32] Haitao Yuan and Guoliang Li. 2021. A Survey of Traffic Prediction: from Spatio-Temporal Data to Intelligent Transportation. *Data Science and Engineering* 6, 1 (2021), 63–85. <https://doi.org/10.1007/s41019-020-00151-z>
- [33] George K. Zipf. 1946. The $P_1 P_2 / D$ Hypothesis: On the Intercity Movement of Persons. *American Sociological Review* 11, 6 (1946), 677–686.

A Implementation Details

A.1 Summary of Overall Approach

Input: GPS trajectory data with ID, latitude, longitude, and timestamp observations

- (1) **Data Preprocessing:** encode trajectory data into occupancy counts following Figure 3. Result: occupancy counts y_t^i for every spatial cell $i \in 1, \dots, I$ and time bin $t \in 1, \dots, T$.
- (2) **Initialize DGLMs:** For all spatial cells $i \in 1, \dots, I$:
 - (a) Select the first \tilde{t} timesteps for initialization (in results above, $\tilde{t} = 288$, or initialization over the first 72 hours with 15 minute temporal resolution)
 - (b) Select the observation distribution based on the mean of the first \tilde{t} occupancy counts, $\bar{y}_{1:\tilde{t}}^i = \frac{1}{\tilde{t}} \sum_{j=1}^{\tilde{t}} y_j^i$, and the sparsity of these occupancy counts $s_{1:\tilde{t}}^i = \frac{1}{\tilde{t}} \sum_{j=1}^{\tilde{t}} \mathbf{1}(y_j^i = 0)$:
 - If $\bar{y}_{1:\tilde{t}}^i > 50$ and $s_{1:\tilde{t}}^i < 0.15$, select a DLM.
 - If $\bar{y}_{1:\tilde{t}}^i \leq 50$ and $s_{1:\tilde{t}}^i < 0.15$, select a Poisson DGLM.
 - If $\bar{y}_{1:\tilde{t}}^i > 50$ and $s_{1:\tilde{t}}^i \geq 0.15$, select a DLMM.
 - If $\bar{y}_{1:\tilde{t}}^i \leq 50$ and $s_{1:\tilde{t}}^i \geq 0.15$, select a DCMM.
- (c) Set the covariate vector F_t and state matrix G_t based on the trend, covariate, and seasonal terms in the model; these values are known ahead of time and F_t and G_t are not learned parameters. In this work, we use a linear trend term, no covariates, and daily seasonality with the first two harmonics (with 15 minute time bins, the period of the seasonality, p , is 96), $F_t = F$, $G_t = G$ and:

$$F = \begin{bmatrix} 1 & 1 & 0 & 1 & 0 \end{bmatrix}',$$

$$G = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \cos(2\pi/p) & \sin(2\pi/p) & 0 & 0 \\ 0 & -\sin(2\pi/p) & \cos(2\pi/p) & 0 & 0 \\ 0 & 0 & 0 & \cos(4\pi/p) & \sin(4\pi/p) \\ 0 & 0 & 0 & -\sin(4\pi/p) & \cos(4\pi/p) \end{bmatrix},$$

- (d) Set and/or tune the discount factors. To tune the discount factors:
- Select a grid of possible values for the discount factor(s) to be tuned.
 - For each value in the grid, initialize the appropriate DGLM with that discount factor value and update the model over all timesteps $t \in 1, \dots, \tilde{t}$. Save the one-step ahead forecasts and calculate the MAE.
 - Select the discount factor over the grid resulting in the lowest MAE.
 - We set the seasonality discount factor to be 0.994, the random effect discount factor or the stochastic variance discount factor to 0.9, and tune the trend discount factor over the grid of values [0.96, 0.97, 1.0].
- (e) Result: an initialized DGLM for each spatial cell, \mathcal{M}^i .
- (3) **Modeling with DGLMs:** In parallel across all spatial cells i and sequentially for $t \in 1, \dots, T$:
- Forecast k steps ahead with \mathcal{M}^i based on information up to time $t - 1$.
 - Update model \mathcal{M}^i using occupancy count y_t^i .
- (4) **Analysis:** Analyze results and calculate metrics.
- Steps (2) and (3) are implemented in PyBats [16].

A.2 DGLM Sequential Learning

DGLMs are Bayesian state space models (Equation 1), and learning proceeds sequentially through time. Below are additional details for step (3b) in the implementation procedure in Section A.1 above for a Poisson DGLM. Sequential analysis is analogous for the other DGLM forms, though the specific forms in steps (3) - (7) are unique to the specific exponential family form. Procedure following [3, 30]; please see these references and [28] for the general case.

For Poisson DGLM \mathcal{M}^i at time t , $y_t^i \sim \text{Poisson}(\eta_t)$, $\lambda_t = \log \eta_t$:

- State vector posterior from time $t - 1$:
 $(\theta_{t-1}|y_{1:t-1}^i) \sim (\mathbf{m}_{t-1}, C_{t-1})$; with $\mathbb{E}(\theta_{t-1}|y_{1:t-1}^i) = \mathbf{m}_{t-1}$ and $\mathbb{V}(\theta_{t-1}|y_{1:t-1}^i) = C_{t-1}$.
- State vector prior at time t : $(\theta_t|y_{1:t-1}^i) \sim (\mathbf{a}_t, \mathbf{R}_t)$; with $\mathbf{a}_t = \mathbf{G}_t \mathbf{m}_{t-1}$ and $\mathbf{R}_t = \mathbf{G}_t C_{t-1} \mathbf{G}_t' + \mathbf{W}_t$.
- Prior for the natural parameter at time t (variational Bayes step): $(\eta_t|y_{1:t-1}^i) \sim \text{Gamma}(\alpha_t, \beta_t)$.
- Evaluate the prior hyper-parameters α_t and β_t such that:
 $\mathbb{E}(\lambda_t|y_{1:t-1}^i) = f_t = \mathbf{F}_t' \mathbf{a}_t$ and $\mathbb{V}(\lambda_t|y_{1:t-1}^i) = q_t = \mathbf{F}_t' \mathbf{R}_t \mathbf{F}_t$.
For the Poisson DGLM, $f_t = \psi(\alpha_t) - \log \beta_t$, and $q_t = \psi'(\alpha_t)$, where ψ is the digamma function and ψ' is the first derivative of the digamma function. We can use numerical optimization, like Newton-Raphson, to solve for α_t and β_t .
- Forecast y_t^i one-step ahead:
 $p(y_t^i|y_{1:t-1}^i) = \text{Negative Binomial}(\alpha_t, \beta_t/(1 + \beta_t))$.
- Posterior for η_t at time t : $p(\eta_t|y_{1:t}^i) = \text{Gamma}(\alpha_t + y_t^i, \beta_t + 1)$.
- Map back to the linear predictor $\lambda_t = \log \eta_t$: posterior mean $g_t = \mathbb{E}(\lambda_t|y_{1:t}^i) = \psi(\alpha_t + y_t^i) - \log(\beta_t + 1)$ and posterior variance $p_t = \mathbb{V}(\lambda_t|y_{1:t}^i) = \psi'(\alpha_t + y_t^i)$.

- State vector posterior at time t : $(\theta_t|y_{1:t}^i) \sim (\mathbf{m}_t, C_t)$, with posterior mean $\mathbf{m}_t = \mathbf{a}_t + \mathbf{R}_t \mathbf{F}_t (g_t - f_t)/q_t$ and posterior variance $C_t = \mathbf{R}_t - \mathbf{R}_t \mathbf{F}_t \mathbf{F}_t' \mathbf{R}_t' (1 - p_t/q_t)/q_t$.

The initial state vector mean is $\mathbf{m}_0 = \mathbf{0}$ and the initial state vector covariance is $C_0 = \mathbf{I}$, the identity matrix.

A.3 Forecasting

One-step ahead forecast distributions are calculated using step (5) in Section A.2 (and the analogous form for the other DGLMs). The marginal multi-step ahead forecasts are calculated similarly. In the multi-step ahead case, step (2) in Section A.2 is updated by iterating the state vectors forward in time; for k -step ahead forecasts, step (2) becomes: $\mathbf{a}_{t+k} = \mathbf{G}^k \mathbf{m}_{t-1}$, $\mathbf{R}_{t+k} = \mathbf{G}^k C_{t-1} (\mathbf{G}^k)'$, where $\mathbf{G}^k = \mathbf{G}_t \times \mathbf{G}_{t+1} \times \dots \times \mathbf{G}_{t+k}$ is the matrix power. \mathbf{W}_t can optionally be added to \mathbf{R}_t to continue discounting while forecasting. Steps (3) - (5) then proceed as in Section A.2 to result in forecasts for k -steps ahead.

B Additional Forecast Example

We include an additional one-step ahead forecast example for a DCMM compared to the seasonal baseline model (Figure 12). The DGLM is able to adapt to the daily occupancy count spike, even when it occurs slightly less than or more than 24 hours from the previous spike; the seasonal baseline, which predicts the occupancy count at time t as the occupancy count observed 24 hours ago is not able to adjust to noise in the seasonal period. Additionally, the DGLM produces probabilistic forecasts, while the seasonal baseline just produces point forecasts.

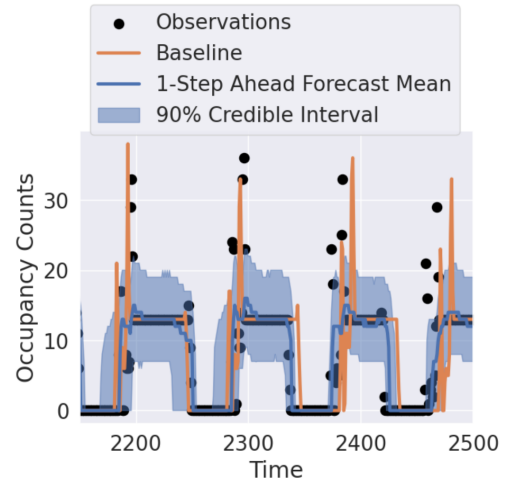


Figure 12: One-step ahead forecast mean and credible intervals, compared to the seasonal baseline, for a DCMM fit to a spatial cell at S2 level 14 with seasonal offset up to 3 hours. The DCMM is better able to adapt to slightly varying seasonal periods.