Enabling the Use of Bayesian Calibration for High-Dimensional Travel Demand Applications via Spatial Input Modeling and the Use of Traffic-Theoretic Priors

Seung Min Baik[1], Giovanni De Nunzio[2], Azise Oumar Diallo[2], Delphine Sinoquet[2], Mohamed Reda El Amri[2], Carolina Osorio[1,*]

[1]Department of Decision Sciences, HEC Montréal, Canada, [2] IFP Energies nouvelles, France

Abstract

Calibrating travel demand in urban traffic simulators is often an underdetermined problem. This leads to multiple OD matrices producing nearly identical simulation outputs, such as traffic counts on certain links of interest. The most popular calibration methods typically aim to estimate a single OD matrix by minimizing discrepancies between simulated and observed traffic counts, but they do not account for the inherent uncertainty in the calibration process. Bayesian calibration (BC) is a general probabilistic framework designed to address this issue by inferring a distribution over plausible calibration parameters rather than identifying a single deterministic value. When applied to travel demand calibration, BC provides a principled approach for uncertainty quantification. However, it struggles with high-dimensional calibration parameters, which is often the case for OD matrices in real-world urban networks, especially due to the computational burden of posterior inference. To alleviate this, we introduce a spatial input modeling approach using a Gaussian mixture model to approximate the OD matrix, transforming the problem dimensionality from the thousands to the dozens while preserving key demand patterns. This allows for a more computationally efficient and interpretable calibration process. We validate our method by estimating the demand for a large-scale case study of the metropolitan area of Lyon, France, using real-world traffic data, demonstrating its ability to quantify uncertainty while maintaining computational efficiency. Our findings underscore the potential of integrating BC with spatial modeling techniques for robust and scalable travel demand calibration in urban transportation systems.

1 Introduction

Accurate calibration of travel demand is a cornerstone of urban transportation planning, playing a key role in various applications, including traffic management, congestion pricing, and infrastructure development. To analyze urban networks, cities are typically divided into multiple traffic analysis zones (TAZs), with each zone serving as either an origin or a destination for trips. The movement of travelers between these zones is represented by an origin-destination (OD) matrix, which quantifies the expected number of trips made between different zone pairs. Since OD matrices are not directly observable, they must be inferred through a calibration process, where the goal is to determine the OD matrix that best aligns with real-world traffic patterns. This is typically done by minimizing the discrepancy between simulated and observed traffic metrics, such as link flows, link speeds, link travel times. As a result, accurate OD matrix calibration is essential for ensuring that transportation models reliably capture network dynamics, leading to data-driven and effective decision-making in urban mobility systems.

This study focuses on offline travel demand calibration, which remains a well-studied yet challenging problem. Several factors contribute to the complexity of calibration, including the high dimensionality of OD matrices, the non-convex relationship between travel demand, which serves as the simulation input, and network performance, and the limited availability of simulated data due to the high computational cost of realistic traffic simulators. To address these issues, various methods have been developed, ranging from general-purpose optimization algorithms, such as stochastic perturbation simultaneous approximation (SPSA) [Lu et al., 2015, Oh et al., 2019] and genetic algorithms [Vaze et al., 2009], to metamodel-based approaches [Osorio, 2019, Dantsuji et al., 2022]. While these approaches achieve successful calibration, they primarily focus on estimating a single OD matrix and do not explicitly consider the uncertainty about the calibrated demand.

Recent studies, thus, have sought to address uncertainties in travel demand calibration, particularly related to the underdetermination issue. These approaches either directly estimate the distribution of OD matrices, for example via variational inference [Mladenov et al., 2022], or quantify the propagated uncertainty associated with travel demand [Wang et al., 2024, Zhuang et al., 2022]. Studying the probabilistic characteristics of travel demand is advantageous as it allows for more robust decision-making in downstream applications that utilize calibrated results. In line

*Corresponding author: carolina.osorio@hec.ca

with these efforts, we adopt a Bayesian calibration (BC) framework, proposed by Kennedy and O'Hagan [2001], to infer a distribution over plausible OD matrices, rather than estimating a single deterministic OD matrix.

Although BC has been widely adopted across various application areas to infer the distribution of calibration parameters [Freni and Mannina, 2010, Hawkins-Daarud et al., 2013, Perrin, 2020], to the best of our knowledge, it has not yet been applied or adapted to transportation problems. This is likely due to it's inability to tackle problems with dimensions higher than roughly a dozen, where the posterior computation requiring multi-dimensional integration becomes increasingly intractable. Such a computational difficulty arises because Bayesian inference expresses the posterior distribution as proportional to the product of the likelihood and prior, without explicitly computing the normalization constant. Consequently, sampling-based methods such as Markov chain Monte Carlo (MCMC) [Andrieu and Thoms, 2008] are commonly employed to approximate the posterior distribution. However, the efficiency of such methods deteriorates significantly in high-dimensional settings, which is often the case for OD matrices in urban mobility networks, making them impractical for travel demand calibration. While various methodologies have been proposed to enhance Bayesian inference for high-dimensional parameters [Chen et al., 2019, Monmarché, 2021], they are usually not well-suited for handling OD matrices as calibration inputs. Unlike generic high-dimensional parameters, OD matrices inherently incorporate spatial relationships, as they represent travel demand between specific zone pairs within a network. Existing general-purpose methods that do not account for such underlying spatial information naturally have limited applicability to travel demand calibration problems.

Therefore, to overcome this challenge, we introduce a spatial input modeling approach to improve computational tractability of posterior inference. Instead of directly calibrating the OD matrix, we define an alternative calibration low-dimensional vector parameter using a Gaussian mixture model (GMM), which provides a continuous mapping from geographical coordinates of an OD pair to its travel demand. Specifically, GMM models the probability density function (PDF) over the coordinate space, where the PDF values represent travel demand. A given set of GMM parameters defines a PDF, from which the OD matrix can be reconstructed by evaluating the density at discrete OD locations. This transformation reduces the dimensionality of the calibration parameter from the scale of thousands (OD matrix) to only dozens (GMM parameters), significantly improving the feasibility of posterior inference while preserving key spatial demand patterns. As a result, BC is performed in this GMM parameter space, enabling a numerically stable and computationally efficient calibration process. In addition, we further improve the sample efficiency by incorporating a physics-informed prior based on traffic theory into the Gaussian process (GP) surrogate model, which is fitted during the calibration. When the amount of simulated data is limited, the physics-informed prior enhances BC's ability to identify well-performing posterior distributions more effectively. We validate our approach on real-world traffic data from Lyon, France, illustrating its scalability and robustness.

Below we summarize the main contributions of this research study:

- We apply BC to offline travel demand calibration, deriving a posterior distribution of the OD matrix instead of a single estimate. This allows for uncertainty quantification, leading to more robust decision-making in urban mobility systems.
- To address the issue arising from the high dimensionality of origin-destination matrix, we propose a spatial input modeling approach using Gaussian mixture models. This allows for an alternative representation of the calibration parameter in a lower-dimensional space, making BC computationally feasible regardless of metropolitan size while preserving key demand structures.
- Additionally, within the BC framework, we integrate physics-informed priors into the Gaussian process surrogate model, leveraging traffic flow theory to enhance calibration accuracy. Our methodology is validated on the urban network of Lyon, France.

The remainder of this abstract is structured as follows. Section 2 describes the BC model for travel demand calibration and the proposed GMM-based spatial input modeling. Next, we outline the contents that will be discussed in detail at the conference, which are omitted here due to page limitations, and then conclude the abstract.

2 Methodology: Enhanced Bayesian Calibration via Spatial Input Modeling

Bayesian Calibration for Travel Demand Calibration For the traffic calibration problem that is of interest here, we consider the following simplified setting of the Kennedy and O'Hagan [2001] model:

$$z = \eta(\boldsymbol{\theta}) + \epsilon, \tag{1}$$

where $\eta(\boldsymbol{\theta})$ represents the discrepancy between simulated and observed traffic counts, quantified as the normalized root mean squared error (NRMSE) $\eta(\boldsymbol{\theta}) = \sqrt{\sum_{s=1}^{S}(c_s - c_s^{sim})^2}/((\sum_{s=1}^{S} c_s)/S)$ with field measurements $\mathbf{c}$ and the simulated counts $\mathbf{c}^{sim}$ across S sensor-equipped links, and ϵ accounts for observational noise. Following the mechanism in Kennedy and O'Hagan [2001], we compute the posterior distribution of the calibration parameter $\boldsymbol{\theta}$ given a target observation z^{tar}, $p(\boldsymbol{\theta}|z = z^{tar})$, using Bayes' theorem, and the corresponding predictive posterior of the new observation.

Spatial Input Modeling with Gaussian Mixture Models BC posterior of the calibration parameter $\boldsymbol{\theta}$ is expressed as a proportionality-based equation rather than an exact PDF, and sampling-based methods such as MCMC are typically used to estimate the posterior distribution. However, due to the curse of dimensionality, direct posterior computation becomes numerically intractable when the OD matrix $\boldsymbol{\theta}$ has thousands of dimensions. To address this, we propose a spatial input modeling approach that approximately represents the OD matrix in a lower-dimensional alternative parameter space, making the posterior inference computationally feasible within the BC framework.

Let the (i, j)-th entry of the OD matrix $\boldsymbol{\theta}$, denoted as $\theta_{i,j}$, represent the travel demand from origin O_i to destination D_j. We define a continuous spatial domain $\mathcal{X}$, where each point $\mathbf{x} = (x_{O,1}, x_{O,2}, x_{D,1}, x_{D,2})$ corresponds to an OD pair, with its origin located at $(x_{O,1}, x_{O,2})$ and its destination at $(x_{D,1}, x_{D,2})$. Since locations in urban networks are typically represented using latitude-longitude coordinates, $\mathcal{X}$ serves as a continuous representation of locations with arbitrary origins and destinations, whereas $\boldsymbol{\theta}$ is a discrete matrix defined over fixed origins and destinations.

To model travel demand over this continuous space, we introduce a spatial function $f : \mathcal{X} \to \mathbb{R}$, which maps an OD pair's coordinates to its corresponding demand value. We assume f is a parametric function characterized by parameters α. Given a specific α, the function $f(\cdot; \alpha)$ fully determines the OD matrix $\boldsymbol{\theta}$ via

$$\theta_{i,j} = f(x_{O_i,1}, x_{O_i,2}, x_{D_j,1}, x_{D_j,2}; \alpha), \tag{2}$$

where $(x_{O_i,1}, x_{O_i,2})$ and $(x_{D_j,1}, x_{D_j,2})$ are the coordinates of the origin and destination, respectively. This formulation implies that optimizing over OD matrices $\boldsymbol{\theta}$ is equivalent to optimizing over α. If we denote the domain of $\boldsymbol{\theta}$ as Θ and that of α as $\mathcal{A}$, our calibration problem is transformed into an exploration over $\mathcal{A}$—a significantly lower-dimensional space compared to Θ.

For modeling the spatial function f, we employ a GMM. Specifically, we define $f(\mathbf{x})$ as the PDF of a GMM with K components multiplied by a scaling factor C as follows:

$$f(\mathbf{x}) := C \sum_{k=1}^{K} \pi_k \mathcal{N}\left(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\right) = C \sum_{k=1}^{N} \pi_k \frac{1}{\sqrt{(2\pi)^4 |\boldsymbol{\Sigma}_k|}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{x} - \boldsymbol{\mu}_k)\right). \tag{3}$$

where π_k are the mixing coefficients, satisfying $\sum_{k=1}^{K} \pi_k = 1$ and $\pi_k \geq 0$, and $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ denotes a PDF of multivariate normal distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$. The full parameter α consists of $\alpha = (C, \pi_1, \ldots, \pi_K, \boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_K, \boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_K)$. For each Gaussian component, the total number of parameters is 15, comprising of a one mixing coefficient, a 4-dimensional mean vector, and a 4×4 symmetric covariance matrix (which has 10 independent elements). Thus, additionally including a one scaling factor, the overall dimensionality of α is $15K + 1$, which remains significantly lower than that of the OD matrix. For example, with $K = 5$, we obtain only 76 parameters.

Bayesian Calibration on a Reduced Parameter Space We now describe how BC can be adapted to the reduced parameter space $\mathcal{A}$ using the GMM representation. Consider a scenario where the OD matrices in the dataset $\mathcal{D}$ were originally generated from a set of GMM parameters $\{\alpha^{(n)}\}_{n=1}^{N}$. In this case, the GMM parameters were first sampled, their PDFs were derived, and then OD matrices were constructed by evaluating the density function at discrete OD locations. Let $\mathbf{X}^{OD} = [\mathbf{x}_{i,j}] = [(x_{O_i,1}, x_{O_i,2}, x_{D_j,1}, x_{D_j,2})]$ be the matrix of coordinate values for the OD pairs of interest. Given this notation, the OD matrix can be expressed as:

$$\boldsymbol{\theta}^{(n)} = f(\mathbf{X}^{OD}; \alpha^{(n)}),\ 1 \leq n \leq N.$$

Alternatively, in cases where OD matrices were originally generated independently and used to run simulations, we can infer the corresponding GMM parameters and scaling factor C by fitting $f(\mathbf{X}^{OD})$ to the OD matrices using methods such as the expectation-maximization algorithm. This process reconstructs an approximate spatial input model, allowing us to express the original OD matrix in terms of GMM parameters $\alpha^{(n)}$. However, the accuracy of this approximation depends on the chosen number of Gaussian components K. A lower K may result in a more compact representation but could fail to capture fine-grained demand patterns, leading to potential loss of calibration accuracy.

In either case, we construct an alternative simulation dataset $\mathcal{E}$ in terms of α as:

$$\mathcal{E} = \{(\alpha^{(n)}, \xi(\alpha^{(n)}))\}_{n=1}^{N}, \quad \text{where} \quad \xi(\alpha^{(n)}) = \eta(f(\mathbf{X}^{OD}; \alpha^{(n)})).$$

Due to the functional relationship between $\alpha^{(n)}$ and $\boldsymbol{\theta}^{(n)}$, the datasets $\mathcal{D}$ and $\mathcal{E}$ are equivalent in terms of their simulation responses. Our approach leverages this equivalence by performing BC, and also the surrogate modeling within it, directly on $\mathcal{E}$ instead of $\mathcal{D}$. Since the dimensionality of α (i.e., $15K + 1$) is significantly lower than that of $\boldsymbol{\theta}$, this transformation enhances computational efficiency while preserving the essential spatial property of travel demand.

Thus, BC can be conducted on $\mathcal{E}$ following the same procedure outlined in Kennedy and O'Hagan [2001]. The only modification required is to replace the observation model $z = \eta(\boldsymbol{\theta}) + \epsilon$ with $z = \xi(\alpha) + \epsilon$. Similarly, the prior and likelihood must be expressed in terms of α rather than $\boldsymbol{\theta}$. However, since α and $\boldsymbol{\theta}$ are functionally linked through $f(\mathbf{X}^{OD}; \alpha)$, this transformation does not impose additional constraints on model specification.

Having the GP-surrogate model fitted on $\mathcal{E}$, denoted as $\xi(\alpha|\mathcal{E})$, the likelihood function of a target observation z^{tar} is expressed in terms of α as $p(z = z^{tar}|\alpha) = \mathbb{P}\left(\xi(\alpha|\mathcal{E}) + \epsilon = z^{tar}\right)$. The posterior distribution of α achieving the minimum discrepancy between the simulated and observed counts (i.e., NRMSE z being zero) is then obtained as:

$$p(\alpha|z = 0) \propto p(z = 0|\alpha) \cdot p(\alpha), \tag{4}$$

via Bayes' theorem. Once the posterior distribution $p(\alpha|z = 0)$ is inferred, the corresponding posterior of the OD matrix $\boldsymbol{\theta}$ can be reconstructed by evaluating $f(\mathbf{X}^{OD}; \alpha)$ for sampled values of α.

Physics-Informed Prior for Gaussian Process In the conference, we will present the formulation for an informative prior distribution that is based on traffic flow theory. It serves the purpose to further enhance sample efficiency.

Anticipated Experimental Results In the conference, we will present numerical experiments conducted on real-world traffic data from Lyon, France, comparing conventional single-point estimation methods with the proposed Bayesian calibration (BC) framework. The results will highlight the impact of spatial input modeling and physics-informed priors on calibration accuracy, sampling efficiency, and uncertainty quantification. Our findings are expected to demonstrate that the proposed approach enhances both the consistency and reliability of OD matrix estimation.

References

Christophe Andrieu and Johannes Thoms. A tutorial on adaptive MCMC. *Statistics and computing*, 18:343–373, 2008.

Peng Chen, Keyi Wu, Joshua Chen, Tom O'Leary-Roseberry, and Omar Ghattas. Projected stein variational newton: A fast and scalable bayesian inference method in high dimensions. *Advances in Neural Information Processing Systems*, 32, 2019.

Takao Dantsuji, Nam H Hoang, Nan Zheng, and Hai L Vu. A novel metamodel-based framework for large-scale dynamic origin–destination demand calibration. *Transportation Research Part C: Emerging Technologies*, 136:103545, 2022.

Gabriele Freni and Giorgio Mannina. Bayesian approach for uncertainty quantification in water quality modelling: The influence of prior distribution. *Journal of Hydrology*, 392(1-2):31–39, 2010.

Andrea Hawkins-Daarud, Serge Prudhomme, Kristoffer G van der Zee, and J Tinsley Oden. Bayesian calibration, validation, and uncertainty quantification of diffuse interface models of tumor growth. *Journal of mathematical biology*, 67(6):1457–1485, 2013.

Marc C Kennedy and Anthony O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.

Lu Lu, Yan Xu, Constantinos Antoniou, and Moshe Ben-Akiva. An enhanced SPSA algorithm for the calibration of dynamic traffic assignment models. *Transportation Research Part C: Emerging Technologies*, 51:149–166, 2015.

Martin Mladenov, Sanjay Ganapathy Subramaniam, Chih wei Hsu, Neha Arora, Andrew Tomkins, Craig Boutilier, and Carolina Osorio. An adversarial variational inference approach for travel demand calibration of urban traffic simulators. In *Proceedings of the 30th ACM SIGSPATIAL Intl. Conf. on Advances in Geographic Information Systems (SIGSPATIAL-22)*, Seattle, WA, 2022.

Pierre Monmarché. High-dimensional MCMC with a standard splitting scheme for the underdamped langevin diffusion. *Electronic Journal of Statistics*, 15(2):4117–4166, 2021.

Simon Oh, Ravi Seshadri, Carlos Lima Azevedo, and Moshe E Ben-Akiva. Demand calibration of multimodal microscopic traffic simulation using weighted discrete SPSA. *Transportation Research Record*, 2673(5):503–514, 2019.

Carolina Osorio. High-dimensional offline origin-destination (OD) demand calibration for stochastic traffic simulators of large-scale road networks. *Transportation Research Part B: Methodological*, 124:18–43, 2019.

Guillaume Perrin. Adaptive calibration of a computer code with time-series output. *Reliability engineering & system safety*, 196:106728, 2020.

Vikrant Vaze, Constantinos Antoniou, Yang Wen, and Moshe Ben-Akiva. Calibration of dynamic traffic assignment models with point-to-point traffic surveillance. *Transportation Research Record*, 2090(1):1–9, 2009.

Qingyi Wang, Shenhao Wang, Dingyi Zhuang, Haris Koutsopoulos, and Jinhua Zhao. Uncertainty quantification of spatiotemporal travel demand with probabilistic graph neural networks. *IEEE Transactions on Intelligent Transportation Systems*, 2024.

Dingyi Zhuang, Shenhao Wang, Haris Koutsopoulos, and Jinhua Zhao. Uncertainty quantification of sparse travel demand prediction with spatial-temporal graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4639–4647, 2022.