

Bayesian Time Varying Coefficient Model with Applications to Marketing Mix Modeling

Edwin Ng
Uber Technologies, Inc
edwinng@uber.com

Zhishi Wang
Uber Technologies, Inc
zhishiw@uber.com

Athena Dai
Uber Technologies, Inc
athena.dai@uber.com

ABSTRACT

Both Bayesian and varying coefficient models are very useful tools in practice as they can be used to model parameter heterogeneity in a generalizable way. Motivated by the need of enhancing Marketing Mix Modeling (MMM) at Uber, we propose a Bayesian Time Varying Coefficient (BTVC) model, equipped with a hierarchical Bayesian structure. This model is different from other time-varying coefficient models in the sense that the coefficients are weighted over a set of local latent variables following certain probabilistic distributions. Stochastic Variational Inference (SVI) is used to approximate the posteriors of latent variables and dynamic coefficients. The proposed model also helps address many challenges faced by traditional MMM approaches. We used simulations as well as real-world marketing datasets to demonstrate our model's superior performance in terms of both accuracy and interpretability.

CCS CONCEPTS

• **Theory of computation** → **Probabilistic computation**; • **Information systems** → **Computational advertising**; • **Mathematics of computing** → **Bayesian computation**.

KEYWORDS

Marketing Mix Modeling, Time Varying Coefficient Model, Hierarchical Bayesian Model, Bayesian Time Series

ACM Reference Format:

Edwin Ng, Zhishi Wang, and Athena Dai. 2021. Bayesian Time Varying Coefficient Model with Applications to Marketing Mix Modeling. In *Proceedings of The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, August 14–18, 2021, Singapore (KDD '21)*. ACM, New York, NY, USA, 6 pages.

1 INTRODUCTION

Marketing as an essential growth driver accounts for sizable investment levels at many companies. Given these large investments, it is not surprising that understanding the return and optimizing the allocation of marketing investment is of foundational importance to marketing practitioners. For many decades Marketing Mix Model (a.k.a. MMM) has been leveraged as one of the most important tools in marketers' arsenal to address such needs. Recent consumer

privacy initiatives (e.g., Apple's announcement of no-IDFA¹ in iOS 14) further underscores the strategic importance of future-proofing any marketing measurement game plan with MMM.

While randomized experiments [26] and causal models [14] are often used for causal inference, they can be either costly or simply infeasible [17] under some circumstances. As an alternative, MMM offers a solution by leveraging aggregated time-series data and regression to quantify the relationship between marketing and demand [19]. MMM could be further tailored and enhanced for different requirements and purposes such as controlling for seasonality, trend, and other control factors [16] and introducing the geo-level hierarchy [24]. More importantly, the primary use case of MMM is often not to predict the sales, but rather to quantify the marginal effects of the different marketing tactics.

There are various issues and challenges that have to be accounted for when building MMM [5]. First, advertising media is evolving at a fast pace so it requires modelers to take into account new marketing levers constantly. It ends up forcing modelers to face the "small n large p " problem. Second, in order to have actionable insights, modelers tend to pick a high level of data granularity. However, higher level of data granularity may lead to sparse observations and outliers. Practitioners need to strike a balance between the limited amount of reliable historical data and a proper level of data granularity. Third, the sequential nature of the data makes it more susceptible to correlated errors, which violates the basic model assumption of ordinary least squares [21]. Fourth, there are severe endogeneity and multicollinearity concerns due to common marketing planning practices and media dynamics. For instance, setting marketing budget as percent of expected revenue is widely used, which contributes to both endogeneity as well as multicollinearity (i.e., highly correlated channel-level spend) in the models. Self-selection bias especially for the demand-capturing channels such as branded paid search [2] can also lead to inflated measurement results if not properly addressed. Fifth, in practice MMM usually involves a large amount of investment and a diverse set of stakeholders with whom alignments need to be secured. As such the bar for model interpretability is very high. Lastly, it is often hard to rely on traditional machine learning approaches such as cross-validation when tuning parameters and choosing models for MMM. Modelers cannot just fall back on that, since there is rarely enough data and/or the holdout periods may not be representative of the challenging series to forecast.

It has been a long journey to build an in-house MMM solution from zero to one at Uber, which takes collaborative efforts across marketers, engineers and data scientists. Throughout this journey, we seek to address all of the above challenges. The preferred modeling solution needs to have the capability of deriving

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '21, August 14–18, Singapore

© 2021 Association for Computing Machinery.

¹<https://developer.apple.com/app-store/user-privacy-and-data-use/>

time-varying elasticity along with other temporal patterns on observational studies. More importantly, randomized experimentation results, which are generally deemed as the golden standard in measuring causality, should be incorporated to calibrate the marginal effects in marketing levers. Benefits from both experimentation and regression modeling can be maximized when combined into one holistic framework.

In this paper, we introduce a class of Bayesian time varying coefficient (BTVC) models that power Uber's MMM solution. Our work brings the ideas of Bayesian modeling and kernel regression together. The Bayesian framework allows a natural way to incorporate experimentation results, and understand the uncertainty of measurement for different marketing levers. The kernel regression is used to produce the time-varying coefficients to capture the dynamics of marketing levers in an efficient and robust way.

The remainder of this paper is organized as follows. In section 2, we describe the problem formulation and the related work. In section 3, we discuss the proposed modeling framework with emphasis on applications to MMM. In section 4, simulations as well as real-case benchmark studies are presented. In section 5, we talk about how to deploy the proposed models using Uber's modern machine learning platform. section 6 is about the conclusion.

2 PROBLEM FORMULATION

2.1 Basic Marketing Mix Model

Expressing sales as a function of spending variables with diminishing marginal returns [9] is one of the fundamental properties in an attribution or marketing response model. In view of that, our model can be expressed in a multiplicative format as below

$$\hat{y}_t = g(t) \cdot \prod_{p=1}^P f_{t,p}(x_{t,p}), \quad t = 1, \dots, T, \quad (1)$$

where $x_{t,p}$ are the regressors (i.e., the ads spending variables in our case), \hat{y}_t is the marketing response, g is a time-series process, f is the cost curve function, P is the number of regressors, and T is the number of time points. Choice of f is desired to have the following properties such that

- \hat{y}_t has an explainable structure to decompose into different driving factors,
- temporal effects such as trend and seasonality of \hat{y}_t are captured,
- $f_{t,p}$ is differentiable and monotonically increasing,
- \hat{y}_t has diminishing marginal returns with respect to $x_{t,p}$.

Equation 1 has an intuitive form as

$$\hat{y}_t = e^{l_t} \cdot e^{s_t} \cdot \prod_{p=1}^P x_{t,p}^{\beta_{t,p}}, \quad 0 \leq \beta_{t,p} \leq 1, \quad \forall t, p, \quad (2)$$

where e^{l_t} is the trend component, e^{s_t} is the seasonality, and $\beta_{t,p}$ are channel-specific time-varying coefficients.

2.2 Related Work

With a log-log transformation, equation (2) can be re-written as

$$\begin{aligned} \ln(\hat{y}_t) &= l_t + s_t + \sum_{p=1}^P \ln(x_{t,p}) \beta_{t,p} \\ &= l_t + s_t + r_t, \quad t = 1, \dots, T, \end{aligned} \quad (3)$$

A natural idea is to use state-space models such as Dynamic Linear Model (DLM) [27] or Kalman filter [7] to solve equation (3). However, there are some caveats associated with these approaches, especially given the goal we want to achieve with MMM:

- DLM with Markov Chain Monte Carlo (MCMC) sampling is not efficient and could be costly especially for high-dimensional problems, which requires sampling for a large number of regressors and time steps.
- Although the Kalman filter provides analytical solutions, it heavily relies on the assumption that the noise is normally distributed, which might be violated in real applications. It also has limited room for further customization such as ingesting priors on coefficients and applying restrictions on coefficient signs (e.g., positive coefficient sign for marketing spend).

Meanwhile, there have been parametric and non-parametric statistical methods proposed [8]. Wu and Chiang [28] considered a nonparametric varying coefficient regression model with longitudinal dependent variable and cross-sectional covariates. Two kernel estimators based on componentwise local least squares criteria were proposed to estimate the time varying coefficients. Li et al. [18] proposed a semiparametric smooth coefficient model as a useful yet flexible specification for studying a general regression relationship with time varying coefficients. It used a local least squares method with a kernel weight function to estimate the smooth coefficient function. Nonetheless, the frequentist approaches can be expensive when doing local estimates with respect to time dimension. There are also no straightforward ways to incorporate information from experimentation results. As such, we are motivated to develop a new approach to derive time varying coefficients under a Bayesian framework for our MMM applications.

3 METHODS

3.1 Time Varying Coefficient Regression

In view of the increased complexity of regression problem in practical MMM, we propose a Bayesian Time Varying Coefficient (BTVC) model, as inspired by the Generalized Additive Models (GAM) [11] and kernel regression smoothing. The key idea behind BTVC is to express regression coefficients as a weighted sum of local latent variables.

First, we define a latent variable $b_{j,p}$ for the p -th regressor at time t_j , $p = 1, \dots, P$, $j = 1, \dots, J$, $t_j \in \{1, \dots, T\}$. There are J latent variables in total for each regressor. From the perspective of spline regression, $b_{j,p}$ can be viewed as a knot distributed at time t_j for a regressor. w is a time-based weighting function such that

$$\beta_{t,p} = \sum_j w_j(t) \cdot b_{j,p}, \quad (4)$$

It is intuitive to use a weighting function taking into account the time distance between t_j and t ,

$$w_j(t) = k(t, t_j) / \sum_{j=1}^J k(t, t_j), \quad (5)$$

where $k(\cdot, \cdot)$ is the kernel function, and the denominator is to normalize the weights across knots. In practice, we have different choices for the kernel functions, such as Gaussian kernel, quadratic kernel or any other custom kernels. In subsection 3.3, we will discuss this in more detail.

We can also rewrite Equation 4 into a matrix form

$$\beta = Kb, \quad (6)$$

where β is the $T \times P$ coefficient matrix with entries $\beta_{t,j}$, K is the $T \times J$ kernel matrix with normalized weight entries $w_j(t)$, and b is the $J \times P$ knot matrix with entries $b_{j,p}$. At time point t , the regression component

$$r_t = X_t \beta_t^T, \quad (7)$$

where $\beta_t = (\beta_{t,1}, \dots, \beta_{t,p})$ and X_t is the t -th row of regressor covariate matrix.

Besides the regression component, we can also apply Equation 4 to other components such as trend and seasonality in Equation 3. Specifically, for the trend component,

$$\begin{aligned} \beta_{lev} &= K_{lev} b_{lev}, \\ l_t &= \beta_{t,lev}. \end{aligned} \quad (8)$$

The trend component can be viewed as a dynamic intercept. For the seasonality component,

$$\begin{aligned} \beta_{seas} &= K_{seas} b_{seas}, \\ s_t &= X_{t,seas} \beta_{t,seas}^T. \end{aligned} \quad (9)$$

$X_{t,seas}$ is the t -th row of seasonality covariate matrix derived from Fourier series. In subsection 3.4, we will discuss the seasonality in more detail.

Instead of estimating the local knots directly (i.e., b , b_{lev} , and b_{seas} in the above equations) by optimizing an objective function, we introduce the Bayesian framework along with customizable priors to conduct the posterior sampling.

3.2 Bayesian Framework

To capture the sequential dynamics and cyclical patterns, we use the Laplace prior to model adjacent knots

$$\begin{aligned} b_{j,lev} &\sim \text{Laplace}(b_{j-1,lev}, \sigma_{lev}), \\ b_{j,seas} &\sim \text{Laplace}(b_{j-1,seas}, \sigma_{seas}). \end{aligned} \quad (10)$$

The initial values ($b_{0,lev}$ and $b_{0,seas}$) can be sampled from a Laplace distribution with mean 0. A similar approach can be found in models implemented in Facebook's Prophet package [25], which uses Laplace prior to model adjacent change points of the trend component.

For the regression component, we introduce a two-layer hierarchy for more robust sampling due to the sparsity in the channel

spending data,

$$\begin{aligned} \mu_{reg} &\sim \mathcal{N}^+(\mu_{pool}, \sigma_{pool}^2), \\ b_{reg} &\sim \mathcal{N}^+(\mu_{reg}, \sigma_{reg}^2), \end{aligned} \quad (11)$$

where the superscript $+$ means a folded normal distribution (positive restriction on the coefficient signs).

In the hierarchy, the latent variable μ_{reg} depicts the overall mean of a set of knots of a single marketing lever. We can treat this as the overall estimate of a channel coefficient across time. This provides two favorable behaviors for the model: during a period with absence of spending for a channel, coefficient knot estimation of such a channel exhibits a shrinkage effect towards the overall estimate; it helps fight against over-fitting due to the volatile local structure.

The two-layer hierarchy is widely adopted in hierarchical Bayesian models and the shrinkage property is sometimes called the pooling effect on regression coefficients [10]. Figure 1 depicts the model flowchart of BTVC. Stochastic Variational Inference [13] is used to estimate the knot coefficient posteriors from which time varying coefficient estimates can be derived using the formulas above in Equation 7, Equation 8, and Equation 9.

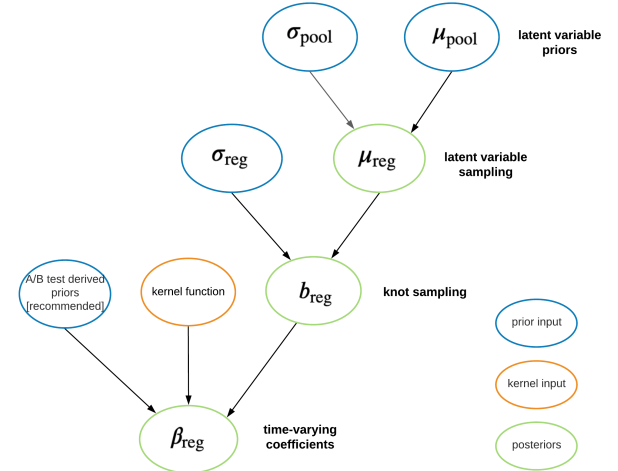


Figure 1: BTVC model flowchart. Blue boxes present the prior related input, where priors derived from lift tests for specific channels can be readily ingested in the framework. Orange box represents the kernel function in use. Green boxes represent the sampled posteriors, which are also the quantities of interest.

3.3 Kernel Selection

For the kernel function used for trend and seasonality, we propose a customized kernel, i.e.,

if $t_j \leq t \leq t_{j+1}$ and $l \in \{j, j+1\}$,

$$k_{lev}(t, t_l) = 1 - \frac{|t - t_l|}{t_{j+1} - t_j}; \quad (12)$$

otherwise zero values are assigned. This kernel bears some similarity with the triangular kernel.

Model	$\beta_1(t)$	$\beta_2(t)$	$\beta_3(t)$
BSTS	0.0067	0.0078	0.0080
tvReg	0.0103	0.0103	0.0096
BTVC	0.0030	0.0026	0.0029

Table 1: Average of mean squared errors based on 100 times simulations.

For the kernel function used for regression, we adopt the Gaussian kernel, i.e.,

$$k_{\text{reg}}(t, t_j; \rho) = \exp\left(-\frac{(t - t_j)^2}{2\rho^2}\right), \quad (13)$$

where ρ is the scale parameter. Other kernels such as Epanechnikov kernel and quadratic kernel etc., could also be leveraged for the regression component.

3.4 Seasonality

Seasonality is a pattern that repeats over a regular period in a time series. To estimate seasonality, a standard approach is to decompose time-series into trend, seasonality and irregular components using Fourier analysis [6]. This method represents the time series by a set of elementary functions called basis such that all functions under study can be written as linear combinations of the elementary functions in the basis. These elementary functions involve the sine and cosine functions or complex exponential. The Fourier series approach describes the fluctuation of time series in terms of sinusoidal behavior at various frequencies.

Specifically, for a given period S and a given order k , two series $\cos(2k\pi t/S)$ and $\sin(2k\pi t/S)$ will be generated to capture the seasonality. For example, with daily data, $S = 7$ represents the weekly seasonality, while $S = 365.25$ represents the yearly seasonality.

4 RESULTS

4.1 Simulations

4.1.1 Coefficient Curve Fitting. We conduct a simulation study based on the following model

$$y_t = \text{trend} + \beta_{1t}x_{1t} + \beta_{2t}x_{2t} + \beta_{3t}x_{3t} + \epsilon_t, \quad t = 1, \dots, T,$$

where the trend and $\beta_{1t}, \beta_{2t}, \beta_{3t}$ are all random walks. The covariates $x_{1t}, x_{2t}, x_{3t} \sim N(3, 1)$ are independent of the error term $\epsilon_t \sim N(0, .3)$.

In our study, we compare BTVC with two other time varying regression models available in R CRAN: Bayesian structural time series (BSTS) [22], and time varying coefficient for single and multi-equation regressions (tvReg) [4]. We set $T=300$ and calculate the average Mean Squared Errors (MSE) against the truth of each regressor across 100 simulations. The estimated coefficient curves of a sample is plotted in Figure 2. The results are reported in Table 1, which demonstrates that BTVC has a better accuracy on the coefficient estimation over the other two models in consideration.

4.1.2 Experimentation Calibration. One appealing property of the BTVC model is its flexibility to ingest any experimentation based priors for any regressors (e.g., advertising channels) since experiments are often deemed as a trustworthy source to tackle the challenges as mentioned in section 1.




Figure 2: Comparison of the BSTS, tvReg and BTVC estimates of coefficient functions. The true values are plotted in grey dots, the blue line is BSTS estimate, the green line refers to tvReg and the red line for BTVC.

To illustrate this feature of BTVC, we first fit a BTVC model on the simulated data, which are generated using a similar simulation scheme as outlined in subsubsection 4.1.1. Next, we assume there is one lift test for the first and third regressors, respectively, and two lift tests for the second regressor. All the tests have a 30-step duration. We use the simulated values as the “truth” derived from the tests, and ingest them as priors into BTVC models. The results are summarized in Figure 3. As expected, the confidence intervals during the ingestion periods and the adjacent neighborhood become narrower, compared to the ones without prior knowledge. Moreover, with this calibration, the coefficient curves are more aligned with the truth around the neighborhood of the test ingestion period. To demonstrate this, in Table 2 we calculated the symmetric mean absolute percentage error (SMAPE) and pinball loss (with 2.5% and 97.5% target quantiles) between the truth and the estimation for the following 30 steps after the prior ingestion period. More importantly, at Uber we apply BTVC to real MMM applications. Its flexibility of ingesting multiple experimentation insights leads to significant improvement in attribution projection accuracy.

4.2 Real Case Studies

To benchmark the model’s forecasting accuracy, we conduct a real-case study using Uber Eats’ data across 10 major countries or markets. Each country series consists of the daily number of first orders in Uber Eats by newly acquired users. The data range spans from Jan 2018 to Jan 2021 including a typical Covid-19 period. The scheme change caused by Covid-19 poses a big challenge for modeling.

We compared BTVC with two other time series modeling techniques, SARIMA [23] and Facebook Prophet [25]. Both Prophet

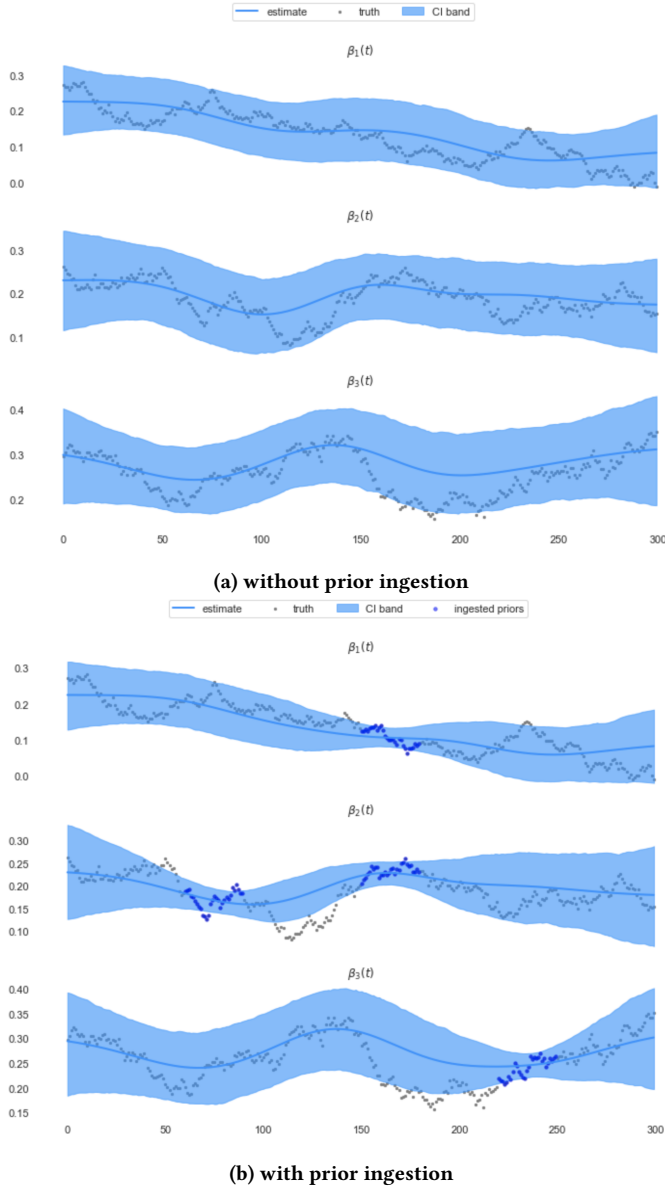


Figure 3: (a) The coefficient estimation of BTVC mode without the prior ingestion. (b) The coefficient estimation of BTVC model with the prior ingestion. The test ingestion periods are highlighted in blue dots. The black solid lines are the simulated truth, while the red ones are the estimations. The shaded bands represent the 95% confidence intervals.

and BTVC models use Maximum A Posterior (MAP) estimates and they are configured as similarly as possible in terms of optimization and seasonality settings. For SARIMA, we fit the $(1, 1, 1) \times (1, 0, 0)_S$ structure by maximum likelihood estimation (MLE) where S represents the choice of seasonality. In our case, $S = 7$ for the weekly seasonality.

	SMAPE		Pinball Loss			
	w/o priors	w priors	w/o priors lower	w/o priors upper	w priors lower	w priors upper
$\beta_1(t)$	0.39	0.21	0.0009	0.0032	0.0005	0.0019
$\beta_2(t)$	1.37	1.25	0.0021	0.0019	0.0011	0.0014
$\beta_3(t)$	0.30	0.18	0.0017	0.0028	0.0014	0.0013

Table 2: SMAPE and pinball loss of coefficient estimates for models without and with prior ingestions. The metrics are calculated using the 30-step coefficients following the test ingestion period. Lower (2.5%) and upper(97.5%) quantiles are reported for pinball loss. With prior ingestion, the coefficient estimation accuracy is improved significantly.

Model	Mean of SMAPE	Std of SMAPE
SARIMA	0.191	0.027
Prophet	0.218	0.067
BTVC	0.174	0.045

Table 3: SMAPE comparison across models. It shows that BTVC outperforms the other two models in terms of average SMAPE across 10 countries.

We use out-of-sample SMAPE as the benchmark performance metric with 28 day forecast horizon on 6 different splits with expanding training windows (i.e., 6 different cuts of the data with incremental training size).

Figure 4 depicts the SMAPE results across the 10 countries, and Table 3 gives the average and standard deviation of SMAPE values for the three models in consideration. As a result, BTVC outperforms the other two models for the majority of 10 countries in terms of SMAPE.




Figure 4: Bar plots of SMAPE on 10 countries with Uber Eats data. Blue bars represent the results from BTVC, orange ones Prophet, and yellow ones SARIMA.

5 ARCHITECTURE

5.1 Implementation

We implemented the BTVC model as a feature branch in our open sourced package *Orbit* [20] by Uber. Orbit is a software package aiming to simplify time series inferences and forecasting with structural Bayesian time series models for real-world cases and research. It provides a familiar and intuitive initialize-fit-predict interface

for time series tasks, while utilizing the probabilistic programming languages such as Stan [3] and Pyro [1] under the hood.

5.2 Deployment

The BTVC deployment system is customized by leveraging Michelangelo [12], a machine learning (ML) platform developed by Uber. Michelangelo provides centralized workflow management for end-to-end modeling process. With the help of Michelangelo, BTVC deployment system is able to automate data preprocessing, model training, validations, predictions, and measurement monitoring at scale.

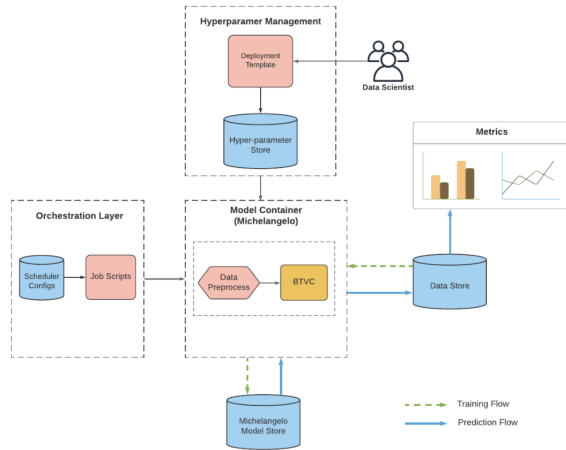


Figure 5: BTVC deployment system.

The deployment workflow summarized in Figure 5 consists of three main components:

- Hyperparameter management layer: this is to store and manage the various supplement data needed for BTVC model training such as normalization scalar, adstock [15], lift test based priors, as well as model specific hyperparameters.
- Orchestration layer: this is to upload and trigger the model training job.
- Model container: a docker container including all the essential modeling code to be integrated with Michelangelo's ecosystem.

6 CONCLUSION

In this paper, we propose a Bayesian Time Varying Coefficient (BTVC) model in particular developed for MMM applications at Uber. By assuming the local latent variables follow certain probabilistic distributions, a kernel-based smoothing technique is applied to produce the dynamic coefficients. This modeling framework entails a comprehensive solution for the challenges faced by traditional MMM. More importantly, it enables marketers to leverage multiple experimentation results in an intuitive yet scientific way. Simulations and real case benchmark studies demonstrate BTVC's superiority in prediction accuracy and flexibility in experimentation ingestion. We also present the model deployment system, which can serve model training and predictions in real time without human oversight or intervention in a scalable way.

7 ACKNOWLEDGMENTS

The authors would like to thank Sharon Shen, Qin Chen, Ruyi Ding, Vincent Pham, and Ariel Jiang for their help on this project, Dirk Beyer, and Kim Larsen for their comments on this paper.

REFERENCES

- [1] Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. 2019. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research* 20, 1 (2019), 973–978.
- [2] Thomas Blake, Chris Nosko, and Steven Tadelis. 2015. Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica* 83, 1 (2015), 155–174.
- [3] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A probabilistic programming language. *Journal of statistical software* 76, 1 (2017).
- [4] Isabel Casas and Ruben Fernandez-Casal. 2021. *tvReg: Time-Varying Coefficients Linear Regression for Single and Multi-Equations*. <https://CRAN.R-project.org/package=tvReg> R package version 0.5.4.
- [5] David Chan and Mike Perry. 2017. Challenges and opportunities in media mix modeling. (2017).
- [6] Alysha M De Livera, Rob J Hyndman, and Ralph D Snyder. 2011. Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American statistical association* 106, 496 (2011), 1513–1527.
- [7] James Durbin and Siem Jan Koopman. 2012. *Time series analysis by state space methods*. Oxford university press.
- [8] Jianqing Fan and Wenyang Zhang. 2008. Statistical methods with varying coefficient models. *Statistics and its Interface* 1, 1 (2008), 179.
- [9] Paul W Farris, Dominique M Hanssens, James D Lenskold, and David J Reibstein. 2015. Marketing return on investment: Seeking clarity for concept and measurement. *Applied Marketing Analytics* 1, 3 (2015), 267–282.
- [10] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. *Bayesian data analysis*. CRC press.
- [11] Trevor J Hastie and Robert J Tibshirani. 1990. *Generalized additive models*. Vol. 43. CRC press.
- [12] Jeremy Hermann and Mike Del Balso. 2017. Meet Michelangelo: Uber's machine learning platform. <https://eng.uber.com/michelangelo/>.
- [13] Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. 2013. Stochastic variational inference. *Journal of Machine Learning Research* 14, 5 (2013).
- [14] Guido W Imbens and Donald B Rubin. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- [15] Yuxue Jin, Yueqing Wang, Yunting Sun, David Chan, and Jim Koehler. 2017. Bayesian methods for media mix modeling with carryover and shape effects. (2017).
- [16] Kim Larsen. 2018. Data Science Can't Replace Human Marketers Just Yet – Here's Why. <https://blog.usejournal.com/data-science-cant-replace-human-marketers-just-yet-here-s-why-d5cdd355d85c>
- [17] Randall A Lewis and Justin M Rao. 2015. The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics* 130, 4 (2015), 1941–1973.
- [18] Qi Li, Cliff J Huang, Dong Li, and Tsu-Tan Fu. 2002. Semiparametric smooth coefficient models. *Journal of Business & Economic Statistics* 20, 3 (2002), 412–422.
- [19] E Jerome McCarthy. 1978. *Basic marketing: a managerial approach*. RD Irwin.
- [20] Edwin Ng, Zhishi Wang, Huigang Chen, Steve Yang, and Slawek Smyl. 2021. Orbit: Probabilistic Forecast with Exponential Smoothing. [arXiv:stat.CO/2004.08492](https://arxiv.org/abs/2004.08492)
- [21] John O Rawlings, Sastry G Pantula, and David A Dickey. 2001. *Applied regression analysis: a research tool*. Springer Science & Business Media.
- [22] Steven L Scott and Hal R Varian. 2014. Predicting the present with bayesian structural time series. *International Journal of Mathematical Modelling and Numerical Optimisation* 5, 1-2 (2014), 4–23.
- [23] Skipper Seabold and Josef Perktold. 2010. statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*. Package version 0.11.1.
- [24] Yunting Sun, Yueqing Wang, Yuxue Jin, David Chan, and Jim Koehler. 2017. Geo-level Bayesian hierarchical media mix modeling. (2017).
- [25] Sean J Taylor and Benjamin Letham. 2018. Forecasting at scale. *The American Statistician* 72, 1 (2018), 37–45. Package version 0.7.1.
- [26] Jon Vaver and Jim Koehler. 2011. Measuring ad effectiveness using geo experiments. (2011).
- [27] Mike West and Jeff Harrison. 2006. *Bayesian forecasting and dynamic models*. Springer Science & Business Media.
- [28] Colin O Wu and Chin-Tsang Chiang. 2000. Kernel smoothing on varying coefficient models with longitudinal dependent variable. *Statistica Sinica* (2000), 433–456.