# Approximate Inference for Non-parametric Bayesian Hawkes Processes and Beyond 

Rui Zhang

# A thesis submitted for the degree of Doctor of Philosophy at The Australian National University 

February 2022
© Copyright by Rui Zhang
All Rights Reserved

Except where otherwise indicated, this thesis is my own original work.

Rui Zhang
28 February 2022

To my parents Zike Zhang, Cuihua Qi, and my fiancée Xiaoyi Chen.

## Chair of Panel

Lexing Xie
Professor of Computer Science, The Australian National University Canberra, ACT, Australia

## Primary Supervisor

Christian Walder
Activity Leader and Senior Researcher, Data61 in CSIRO
Canberra, ACT, Australia

## Associate Supervisor

Marian-Andrei Rizoiu
Lecturer of Computer Science, The University of Technology Sydney Sydney, NSW, Australia

## Acknowledgments

I would like to express my sincere gratitude to the people supporting me during the doctoral journey:

1. My supervisors: Prof. Lexing Xie, Dr. Christian Walder and Dr. Marian-Andrei Rizoiu. It is an extraordinary honor for me to have Lexing, Christian and Andrei to be my advisors. They have immense knowledge in various fields and a strong dedication to research, which motivates me to become a good researcher. They provide patient and continuous guidance and support of my research, which helped this thesis to become a reality.
2. The Research School of Computer Science at ANU and Data61 in CSIRO. I would like to thank both institutes for providing me the PhD scholarships. Their healthy workspace and sufficient resources are crucial for me to develop research skills and explore fields of interest.
3. Members and Alumni of the ANU Computational Media Lab: Dongwoo Kim, Quyu Kong, Swapnil Mishra, Umanga Bista, Siqi Wu, and many others. Much of my work benefits from the discussions with them. I am grateful that I can work with a group of supportive talents.
4. External collaborators: Krikamol Muandet and Bernhard Schölkopf at Max Planck Institute for Intelligent Systems, Masaaki Imaizumi at the University of Tokyo and Edwin Bonilla at Data61 in CSIRO. Their rigorous research spirits have made a profound influence on me.
5. National Computational Infrastructure (NCI). I acknowledge the computational support from NCI which is a very helpful research platform to my PhD .
6. My parents Zike Zhang and Cuihua Qi, my fiancée Xiaoyi Chen and my relatives. I would like to thank my relatives for their support and care. Most importantly, I would like to express deep thanks to my parents Zike Zhang and Cuihua Qi and my fiancée Xiaoyi Chen. They gave me unconditional love and encouragement throughout my PhD.

## Abstract

The Hawkes process has been widely applied to modeling self-exciting events including neuron spikes, earthquakes and tweets. To avoid designing parametric triggering kernels, the non-parametric Hawkes process has been proposed, in which the triggering kernel is in a non-parametric form. However, inference in such models suffers from poor scalability to large-scale datasets and sensitivity to uncertainty in the random finite samples. To deal with these issues, we employ Bayesian non-parametric Hawkes processes and propose two kinds of efficient approximate inference methods based on existing inference techniques. Although having worked as the cornerstone of probabilistic methods based on Gaussian process priors, most of existing inference techniques approximately optimize standard divergence measures such as the Kullback-Leibler (KL) divergence, which lacks the basic desiderata for the task at hand, while chiefly offering merely technical convenience. In order to improve them, we further propose a more advanced Bayesian inference approach based on the Wasserstein distance, which is applicable to a wide range of models. Apart from these works, we also explore a robust frequentist estimation method beyond the Bayesian field. Efficient inference techniques for the Hawkes process will help all the different applications that it already has, from earthquake forecasting, finance to social media. Furthermore, approximate inference techniques proposed in this thesis have the potential to be applied to other models to improve robustness and account for uncertainty.

More specifically, we first develop an efficient non-parametric Bayesian estimation of the triggering kernel of Hawkes processes based on Gibbs sampling. Our method considers a Gaussian process modulated triggering kernel and is developed based on the cluster representation of Hawkes processes. Utilizing the finite support assumption of the Hawkes process, we efficiently sample random branching structures and thus, we split the Hawkes process into clusters of Poisson processes. We derive two algorithms - a block Gibbs sampler and a maximum a posteriori estimator based on expectation maximization - and we show that our methods have a linear time complexity, both theoretically and empirically. On synthetic data, we show our methods to be able to infer flexible Hawkes triggering kernels. On two large-scale Twitter diffusion datasets, we show that our methods outperform the current state-of-the-art in goodness-of-fit and that the time complexity is linear in the size of the dataset. We also observe that on diffusions related to online videos, the learned kernels reflect the perceived longevity for different content types such as music or pet videos.

Secondly, we propose a new non-parametric Bayesian Hawkes process in which the triggering kernel is modeled as a squared sparse Gaussian process and a novel
variational inference schema is proposed for model optimization. We again employ the branching structure of the Hawkes process so that maximization of evidence lower bound (ELBO) is tractable by the expectation-maximization algorithm. We propose a tighter ELBO which improves the fitting performance. Further, we accelerate the novel variational inference scheme to linear time complexity by leveraging the finite support assumption of the triggering kernel. Different from prior acceleration methods, ours enjoys higher efficiency. Finally, we exploit synthetic data and two large social media datasets to evaluate our method. We show that our approach outperforms state-of-the-art non-parametric frequentist and Bayesian methods. We validate the efficiency of our accelerated variational inference scheme and the practical utility of our tighter ELBO for model selection. We observe that the tighter ELBO exceeds the common one in model selection.

Thirdly, we develop a new approximate inference method for Gaussian process models which overcomes the technical challenges arising from abandoning those convenient divergences. Our method-dubbed Quantile Propagation (QP)—is similar to expectation propagation (EP) but minimizes the $L_{2}$ Wasserstein distance instead of the KL divergence. The Wasserstein distance exhibits all the required properties of a distance metric, while respecting the geometry of the underlying sample space. We show that QP matches quantile functions rather than moments as in EP and has the same mean update but a smaller variance update than EP, thereby alleviating EP's tendency to over-estimate posterior variances. Crucially, despite the significant complexity of dealing with the Wasserstein distance, QP has the same favorable locality property as EP, and thereby admits an efficient algorithm. Experiments on classification and Poisson regression show that QP outperforms both EP and variational Bayes.

Finally, we propose a simple and robust framework for the estimation of conditional moment restriction (CMR) models which include the Hawkes process. The framework is developed based on a kernelized CMR known as a maximum moment restriction (MMR) and applied to nonlinear instrumental variable (IV) regression, which we are particularly interested in. The MMR is formulated by maximizing the interaction between the residual and the instruments belonging to a unit ball in a reproducing kernel Hilbert space (RKHS). The MMR allows us to reformulate the IV regression as a single-step empirical risk minimization problem, where the risk depends on the reproducing kernel on the instrument and can be estimated by a Ustatistic or V-statistic. This simplification not only eases the proofs of consistency and asymptotic normality in both parametric and non-parametric settings, but also results in easy-to-use algorithms with an efficient hyper-parameter selection procedure. We demonstrate the advantages of our framework over existing ones using experiments on both synthetic and real-world data.

## Publications and Software

Published and under-reviewed papers are summarized as follows. The software developed in these works is also provided.

- Zhang, R.; Muandet, K.; Sснölкopf, B.; and Imaizumi, M., 2021b. Instrument space selection for kernel maximum moment restriction. arXiv preprint, (2021). Submitted to NeurIPS 2021. Paper: https://arxiv.org/abs/2106.03340. Software: https://github.com/RuiZhang2016/Instrument-Space-Selection-for-Kernel-Maxi mum-Moment-Restriction. (Not included)
- Zhang, R.; Imaizumi, M.; Schölkopf, B.; and Muandet, K., 2021a. Maximum moment restriction for instrumental variable regression. arXiv preprint, (2021). Submitted to NeurIPS 2021. Paper: https://arxiv.org/abs/2010.07684. Software: https://github.com/RuiZhang2016/MMRIV. (In Chapter 6)
- Zhang, R.; Walder, C.; Bonilla, E. V.; Rizoiu, M.-A.; and Xie, L., 2020a. Quantile propagation for wasserstein-approximate gaussian processes. In Advances in Neural Information Processing Systems (NeurIPS 2020), vol. 33, 21566-21578. Curran Associates, Inc. Paper: https://papers.nips.cc/paper/2020/hash/f5e62af8852 93cf4d511ceef31e61c80-Abstract.html. Software: https://github.com/RuiZhan g2016/Quantile-Propagation-for-Wasserstein-Approximate-Gaussian-Processes. (In Chapter 5)
- Zhang, R.; Walder, C.; and Rizoiv, M.-A., 2020b. Variational inference for sparse gaussian process modulated hawkes proces. In The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020). New York, U.S.A. Paper: https: //arxiv.org/abs/1905.10496. Software: https://github.com/RuiZhang2016/Vari ational-Inference-for-SGP-Modulated-Hawkes-Process. (In Chapter 4)
- Zhang, R.; Walder, C.; Rizoiu, M.-A.; and Xie, L., 2019. Efficient non-parametric bayesian hawkes processes. In International Joint Conference on Artificial Intelligence (IJCAI 2019). Macao, China. Paper: http://arxiv.org/abs/1810.03730. Software: https://github.com/RuiZhang2016/Efficient-Nonparametric-Bayesian-Hawkes -Processes. (In Chapter 3)


## Contents

Acknowledgments ..... v
Abstract ..... vii
Publications and Software ..... ix
1 Introduction ..... 1
1.1 Research Questions ..... 2
1.2 Thesis Contributions ..... 3
1.2.1 Laplace Bayesian Hawkes Process ..... 4
1.2.2 Variational Bayesian Hawkes Process ..... 4
1.2.3 Quantile Propagation for Gaussian Process Models ..... 5
1.2.4 Kernel Maximum Moment Restriction Estimation ..... 6
1.3 Broader Impact ..... 6
1.4 Thesis Outline ..... 6
2 Preliminaries and Related Work ..... 7
2.1 Poisson and Hawkes Processes ..... 7
2.1.1 Poisson Processes ..... 7
2.1.2 Hawkes Processes ..... 8
2.2 Gaussian Processes ..... 9
2.2.1 Exact Gaussian Processes ..... 9
2.2.2 Sparse Gaussian Processes ..... 10
2.2.3 Laplace Approximation ..... 11
2.2.4 Variational Inference ..... 12
2.2.5 Expectational Propagation ..... 13
2.3 Wasserstein Distance ..... 14
2.4 Conditional Moment Restriction ..... 15
2.4.1 Instrumental Variable Regression ..... 16
2.5 Related Works of Hawkes Process Estimation ..... 17
2.5.1 Parametric Frequentist Solutions ..... 17
2.5.2 Non-parametric Frequentist Solutions ..... 17
2.5.3 Bayesian Parametric and Non-parametric Solutions ..... 18
2.6 Related Works of Expectation Propagation and Wasserstein Distance ..... 19
2.7 Related Works of Conditional Moment Restriction ..... 20
2.8 Summary ..... 22
3 Gibbs Sampling and Laplace Approximation Based Efficient Inference ..... 23
3.1 Laplace Bayesian Poisson Process ..... 24
3.2 Inference via Sampling ..... 25
3.2.1 Distribution and Sampling of the Branching Structure ..... 25
3.2.2 Posterior Distribution of $\mu$ ..... 26
3.2.3 Posterior Distribution of $\phi$ ..... 26
3.2.4 Computational Complexity ..... 27
3.3 Maximum-A-Posterior Estimation ..... 28
3.3.1 Relationship to EM ..... 28
3.3.2 EM-Hawkes ..... 28
3.4 Experiments ..... 29
3.4.1 Synthetic Data ..... 30
3.4.2 Twitter Diffusion Data ..... 32
3.5 Summary ..... 33
4 Variational Inference for Sparse Gaussian Process Modulated Hawkes Pro- cess ..... 35
4.1 Variational Bayesian Poisson Process ..... 36
4.2 Variational Bayesian Hawkes Process ..... 37
4.2.1 Notations ..... 37
4.2.2 Data Dependent Expectation ..... 38
4.2.3 Predictive Distribution of $\phi$ ..... 39
4.3 New Variational Inference Schema ..... 39
4.3.1 New Optimization Schema for VBHP ..... 41
4.4 Acceleration Trick ..... 42
4.4.1 Time Complexity Without Acceleration ..... 42
4.4.2 Acceleration to Linear Time Complexity ..... 42
4.5 Experiments ..... 42
4.5.1 Evaluation ..... 42
4.5.2 Prediction ..... 43
4.5.3 Baselines ..... 43
4.5.4 Synthetic Experiments ..... 45
4.5.5 Real World Experiments ..... 46
4.6 Conclusions ..... 46
5 Quantile Propagation for Wasserstein-Approximate Gaussian Processes ..... 49
5.1 Introduction ..... 50
5.2 Quantile Propagation ..... 51
5.2.1 Convexity of $L_{p}$ Wasserstein Distance ..... 51
5.2.2 Minimization of $L_{2}$ WD ..... 51
5.2.3 Properties of the Variance Update ..... 52
5.3 Locality Property ..... 53
5.3.1 Review: Locality Property of EP ..... 53
5.3.2 Locality Property of QP ..... 54
5.4 Experiments ..... 54
5.4.1 Binary Classification ..... 55
5.4.2 Poisson Regression ..... 56
5.5 Conclusions ..... 57
6 Kernel Maximum Moment Restriction for Instrumental Variable Regres- sion ..... 59
6.1 Introduction ..... 60
6.2 Preliminaries ..... 60
6.2.1 Maximum Moment Restriction ..... 61
6.3 Our Method ..... 62
6.3.1 Empirical Risk Minimization with U, V-Statistics ..... 63
6.3.2 Practical MMR-IV Algorithms ..... 64
6.4 Hyper-parameter Selection ..... 65
6.5 Consistency and Asymptotic Normality ..... 66
6.5.1 Consistency ..... 67
6.5.2 Asymptotic Normality ..... 67
6.6 Experimental Results ..... 68
6.6.1 Low-dimensional Scenarios ..... 68
6.6.2 High-dimensional Structured Scenarios ..... 69
6.6.3 Mendelian Randomization ..... 70
6.6.4 Application on the Vitamin D data ..... 71
6.7 Conclusion ..... 72
7 Conclusions and Future Work ..... 73
7.1 Conclusions ..... 73
7.2 Future Research Directions ..... 74
A Appendix: Laplace Bayesian Hawkes Process ..... 75
A. 1 Computing the Integral Term of the Log-likelihood ..... 75
A. 2 M.A.P. $\mu$ and $\phi$ Given Infinite Branching Structures ..... 76
A. 3 Mode-Finding the Triggering Kernel ..... 76
B Appendix: Variational Bayesian Hawkes Process ..... 79
B. 1 Deriving Equation (4.5) ..... 79
B. 2 Closed Forms of KL Terms in the ELBO ..... 80
B. 3 Extra Closed Form Expressions for Equation (4.2.2) ..... 81
B. 4 Additional Experiment Results ..... 81
C Appendix: Quantile Propagation ..... 83
C. 1 Minimization of $L_{2}$ WD between Univariate Gaussian and Non-Gaussian Distributions ..... 83
C. 2 Minimization of $L_{p}$ WD between Univariate Gaussian and Non-Gaussian Distributions ..... 84
C. 3 Computations for Different Likelihoods ..... 84
C.3.1 Probit Likelihood for Binary Classification ..... 85
C.3.2 Square Link Function for Poisson Regression ..... 86
C. 4 Proof of Convexity ..... 88
C. 5 Proof of Variance Difference ..... 89
C. 6 Proof of Locality Property ..... 89
C.6.1 Details of Equation 5.2 ..... 91
C.6.2 Details of Equation C. 3 ..... 92
C.6.3 Details of Equation C. 12 ..... 92
C.6.4 Details of Equation C. 12 ..... 93
C. 7 More Details of EP ..... 93
C. 8 Predictive Distributions of Poisson Regression ..... 94
C. 9 Proof of Corollary 3.2 ..... 94
C. 10 Lookup Tables ..... 95
D Appendix: Kernel Maximum Moment Restriction ..... 97
D. 1 Integrally Strictly Positive Definite (ISPD) Kernels ..... 97
D. 2 Detailed Proofs ..... 97
D.2.1 Proof of Lemma 1 ..... 98
D.2.2 Proof of Theorem 6 ..... 98
D.2.3 Convexity Result ..... 98
D.2.4 Uniform Convergence of Risk Functionals ..... 99
D.2.5 Indefiniteness of Weight Matrix $W_{U}$ ..... 101
D.2.6 Consistency of $\hat{f}_{V}$ with Convex $\Omega(f)$ ..... 101
D.2.7 Proof of Theorem 8 ..... 102
D.2.8 Consistency of $\hat{f}_{U}$ ..... 102
D.2.9 Asymptotic Normality of $\hat{\theta}_{U}$ ..... 102
D.2.10 Proof of Theorem 9 ..... 106
D.2.11 Asymptotic Normality in the Infinite-dimension Case ..... 108
D. 3 Gaussian Process (GP) Interpretation ..... 115
D.3.1 Likelihood Function ..... 116
D.3.2 Maximum A Posteriori (MAP) ..... 117
D. 4 Related Works in Reinforcement Learning ..... 118
D. 5 Additional Details of Experiments ..... 120
D.5.1 Baseline Algorithms ..... 120
D.5.2 Experimental Settings on Low-dimensional Scenario ..... 121
D.5.3 Experimental Settings on High-dimensional Scenario ..... 122
D.5.4 More Details of MMR-IVs ..... 123
D.5.5 Additional Comments on Mendelian Randomization ..... 123
D.5.6 Experimental Settings on Vitamin D Data ..... 124

## List of Figures

1.1 The cluster representation of a Hawkes process. In (a), a Hawkes process with a decaying triggering kernel $\phi(\cdot)$ has intensity $\lambda(\boldsymbol{x})$ which increases after each new point (vertical dash line) is generated. It can be viewed as a cluster of Poisson processes: $\operatorname{PP}(\mu)$ and $\operatorname{PP}\left(\phi\left(x-x_{i}\right)\right)$ associated with each $x_{i}$. Figure (b) presents the branching structure of the Hawkes process in (a) and it reflects the triggering relationships between points. Here, an edge $\boldsymbol{x}_{i} \rightarrow \boldsymbol{x}_{j}$ means that $\boldsymbol{x}_{i}$ triggers $\boldsymbol{x}_{j}$, and its probability is denoted as $p_{j i}$. ..... 2
2.1 A Poisson process realization (points). The intensity function $\lambda(t)=$ $10 \sin (t)+10.5$ (solid line) is used to simulate points and higher function values generates more points. ..... 8
2.2 A Hawkes process realization (points). The intensity function $\lambda(t)=1+$ $\sum_{t_{i}<t} 0.1 e^{-\left(t-t_{i}\right)}+0.6 e^{-3\left(t-t_{i}\right)}+0.7 e^{-7\left(t-t_{i}\right)}$ (solid line) is used to simulate points. ..... 9
2.3 A causal graph depicting an instrumental variable $Z$ that satisfies an exclusion restriction and unconfoundedness (there may be a confounder $\varepsilon$ acting on $X$ and $Y$, but it is independent of $Z$ ) ..... 16
3.1 A visual summary of the Gibbs-Hawkes, EM-Hawkes and the EM algo- rithms. The differences between them are (1) the number of sampled branching structures and (2) selected $\phi$ and $\mu$ for $p_{i j}$. In contrast with with Gibbs-Hawkes, the EM-Hawkes method draws multiple branching structures at once and calculates $p_{i j}$ using M.A.P. $\phi$ and $\mu$. The EM algorithm is equivalent to sampling infinite branching structures and exploiting M.A.P. or constrained M.L.E. $\phi$ and $\mu$ to calculate $p_{i j}$ (see Section 3.3). ..... 28
3.2 Computation time (seconds) for calculating $p_{i j}$ and sampling branching structures, with and without Halpin's speed up. ..... 29
3.3 Running time (seconds) per iteration on ACTIVE and SEISMIC. ..... 29
3.4 Learned Hawkes triggering kernels using our non-parametric Bayesian approaches. Each red or blue area shows the estimated posterior distributions of $\phi$, while the solid lines indicate the 10,50 and 90 percentiles. In Figure (a), a synthetic dataset simulated using $\phi_{\exp }(t)$ (in gray) is fit using Gibbs-Hawkes (in red) and EM-Hawkes (in blue); Figure (b) presents learning outcomes on Twitter data in ACTIVE (in red) and SEISMIC (in blue); Figure (c) presents learning outcomes on Twitter data associated with two categories in the ACTIVE set: Music (in red) and Pets \& Animals (in blue).
4.1 The relationship between the log marginal likelihood and the $\mathrm{L}_{2}$ distance. In (a), the true $\phi_{\sin }$ (dash green) is plotted with the median (solid) and the [0.1, 0.9] interval (filled) of the approximate posterior triggering kernel obtained by VBHP and Gibbs Hawkes (10 inducing points). It uses the maximum point of the TELBO (red star in (b)). In (c), the maximum point of the TELBO is marked. The maximum point overlaps with that of the CELBO. $[0,1.4]$ is used as the support of the predictive triggering kernel and 10 inducing points are used.43
4.2 Figure (a), (b): The relationship between the TELBO and the HLL. Figure (c), (d): Average fitting time (seconds) per iteration. In Figure (a), the maximum point is marked by the red star. In Figure (b), the maximum points of the TELBO and CELBO are marked by red and blue stars. Figure (c) is plotted on 50 processes. Figure (d) shows the fitting time of Gibbs Hawkes (star) and VBHP (circle) on 120 processes. 10 inducing points are used unless specified.44
5.1 A scatter plot of the predictive variances of latent functions on test data, for EP and QP. The diagonal dash line represents equivalence. We see that the predictive variance of QP is always less than or equal to that of EP.57
6.1 Runtime comparison in the large-sample regime $(n=2000)$. The computational time of parameter selection is excluded from the comparison. AGMM-K (No Nyström) and MMR-IV (RKHS) overlap due to the same runtime.69
6.2 The MSE of different methods on Mendelian randomization experiments as we vary the numbers of instruments (left), the strength of confounders to exposures $c_{1}$ (middle), and the strength of confounders to instruments $c_{2}$ (right). The MSE is obtained from 10 repetitions of the experiment.70
A. 1 Triggering kernels estimated by the Gibbs-Hawkes method (Section 3.2) and the EM-Hawkes method (Section 3.3.2). The true kernel is plotted as the bold gray curve. We plot the median (red) and [0.1, 0.9] interval (filled red) of the approximate predictive distribution, along with the triggering kernel inferred by the EM Hawkes method (blue). The hyper-parameters $a$ and $b$ of the Gaussian process kernel are set to 0.002. 77
B. 1 Posterior Triggering Kernels Inferred By VBHP and Gibbs Hawkes. Results of Gibbs Hawkes are obtained in 2000 iterations.82
B. 2 Convergence Rate of VBHP and Gibbs Hawkes with Different Numbers of Inducing Points. VBHP and Gibbs Hawkes measure respectively the relative error of the approximate marginal likelihood and of the posterior distribution of the Gaussian process.82
D. 1 Estimated effect of vitamin D level on mortality rate, controlled by age. All plots depict normalized contours of $\hat{f}^{\prime}$ defined in Section D.5.6 where blue represents low mortality rate and yellow the opposite. We can divide each plot roughly into left (young age) and right (old age) parts. While the right parts reflect similar information (i.e., lower vitamin D level at an old age leads to higher mortality rate), the left parts are different. In (a), a high level of vitamin $D$ at a young age can result in a high mortality rate, which is counter-intuitive. A plausible explanation is that it is caused by some unobserved confounders between vitamin D level and mortality rate. In (b), on the other hand, this spurious effect disappears when the filaggrin mutation is used as instrument, i.e., a low vitamin D level at a young age has only a slight effect on death, but a more adverse effect at an old age [Meehan and Penckofer, 2014]. This comparison demonstrates the benefit of an instrument variable. (c) and (d) correspond to the results obtained by using Sjolander and Martinussen [2019]' generalized linear model (GLM), from which we can draw similar conclusions. It is noteworthy that MMRIV allows more flexible non-linearity for causal effect.126
D. 2 Distribution of Vitamin D data. Data points are plotted in the middle, the solid curve and histogram on the right describe the kernel density estimation and histogram of Vitamin D, and those on the top are for Age. 127

## List of Tables

2.1 Non-parametric triggering kernel estimation. ..... 19
3.1 Time complexity. ..... 27
3.2 Empirical performance comparison between algorithms (columns) with different measures (rows). Top: relative L2 distance to known $\phi$ and $\mu$, and AVG denoted the average of L2 errors of $\phi$ and $\mu$. bottom: mean predictive log likelihood on real data. Bold numbers denote the best performance and the underlined numbers for the second best. ..... 30
4.1 Notations. ..... 36
4.2 Results on synthetic and real-world data (mean $\pm$ one standard vari- ance). VBHP (C) and (T) use the CELBO and the TELBO to update the hyper-parameters respectively. Bold numbers denote the best perfor- mance. ..... 44
5.1 Results on benchmark datasets. The first three columns give dataset names, the number of instances $m$ and the number of features $n$. The table records the test errors (TEs) and the negative test log-likelihoods (NTLLs). The top section is on the benchmark datasets employed by Kuss and Rasmussen [2005] and the middle section uses additional datasets. The bottom section shows Poisson regression results. * in- dicates that QP outperforms EP in more than $90 \%$ of experiments consistently. ..... 55
6.1 The mean square error (MSE) $\pm$ one standard deviation in the large- sample regime ( $n=2000$ ). ..... 68
6.2 The mean square error (MSE) $\pm$ one standard deviation on high- dimensional structured data. We run each method 10 times. ..... 71
D. 1 The mean square error (MSE) $\pm$ one standard deviation in the small- sample regime ( $n=200$ ). ..... 122
D. 2 Hyper-parameters of neural networks used in the experiments ..... 124

## Chapter 1

## Introduction

Hawkes (point) processes [Hawkes, 1971] have been widely used to model self-exciting event sequences in which existing events increase the likelihood of occurrence of future events. An example of self-exciting data is that on the online social media, attractive tweets can easily diffuse by being retweeted [Simma and Jordan, 2010], and retweeting promotes these tweets being seen and retweeted by more people. As another example, in the seismic activities, an intensive earthquake tends to trigger a series of aftershocks, and aftershocks possibly cause more aftershocks [Zhao et al., 2015]. In addition, many finance data also have the self-excitement characteristic [Bacry et al., 2015] - we recommend related works therein for more self-exciting cases.

The Hawkes process is suitable for self-exciting events because it explicitly models the self-exciting interactions between events. Specifically, every event in the Hawkes process is assumed to be triggered by either a previous event or the exterior stimulus. To quantify the former triggering factor, the model uses a triggering kernel function $\phi(\cdot)$, and for the latter, it employs a background intensity function $\mu(\cdot) . \phi$ and $\mu$ measure the occurrence rate of events, and are non-negative-valued functions. The aggregate occurrence rate is the sum of $\mu$ and $\phi$ of every previous event, which is usually denoted as a $\lambda(\cdot)$ function. Consequently, we can understand a Hawkes process as a cluster of Poisson processes [Hawkes and Oakes, 1974]. In the cluster view, a Poisson process with an intensity $\mu$ (denoted as $\operatorname{PP}(\mu)$ ) generates immigrant points which arrive in the system from the outside, and every existing point triggers offspring points, which are generated internally through self-excitement, via a $\operatorname{PP}(\phi)$. Points are therefore structured into clusters where each cluster contains either a point and its direct offspring or the background process (an example is shown in Figure 1.1(a)). Connecting all points using the triggering relations yields a tree structure, which is called the branching structure (an example is shown in Figure 1.1(b) corresponding to Figure 1.1(a)). With the branching structure, we can decompose the Hawkes process into a cluster of Poisson processes. The triggering kernel $\phi$ is shared among all cluster Poisson processes relating to a Hawkes process, and it determines the overall behavior of the process. Consequently, designing the kernel functions is of utmost importance for employing the Hawkes process to a new application, and its study has attracted much attention.


Figure 1.1: The cluster representation of a Hawkes process. In (a), a Hawkes process with a decaying triggering kernel $\phi(\cdot)$ has intensity $\lambda(\boldsymbol{x})$ which increases after each new point (vertical dash line) is generated. It can be viewed as a cluster of Poisson processes: $\operatorname{PP}(\mu)$ and $\operatorname{PP}\left(\phi\left(\boldsymbol{x}-\boldsymbol{x}_{i}\right)\right)$ associated with each $\boldsymbol{x}_{i}$. Figure (b) presents the branching structure of the Hawkes process in (a) and it reflects the triggering relationships between points. Here, an edge $x_{i} \rightarrow x_{j}$ means that $\boldsymbol{x}_{i}$ triggers $\boldsymbol{x}_{j}$, and its probability is denoted as $p_{j i}$.

### 1.1 Research Questions

Most of the recent works employing Hawkes Processes [Bao et al., 2015; Filimonov and Sornette, 2015; Lallouache and Challet, 2016; Mishra et al., 2016; Rizoiu et al., 2017] design parametric kernels modeling a predetermined subset of social processes such as the limited length of collective memory [Wu and Huberman, 2007], or preferential attachment [Barabási, 2005]. Manually designing parametric kernels is an expensive process and may not generalize well to other applications. An open question is (Q1) can we design non-parametric solutions for the kernel function?

There are many non-parametric estimations of Hawkes triggering kernels, such as the works of Lewis and Mohler [2011]; Zhou et al. [2013]; Bacry and Muzy [2016]; Eichler et al. [2017]. These are all frequentist methods and among them, the Wiener-Hopf equation based method [Bacry and Muzy, 2016] takes the advantage of quadrature approximation for the integrals in the equation and obtains linear time complexity in estimating the triggering kernel, while it is sensitive to the employed quadrature points. A different class of estimation methods are based on the EulerLagrange equation [Lewis and Mohler, 2011; Zhou et al., 2013]. Similarly, these methods require discretizing the input domain and as a result, they face a problem of poorly scaling with the dimension of the domain. The same problem is also faced by Eichler et al. [2017]'s discretization based method. Besides, frequentist methods do not model uncertainty over the learned triggering kernels and tend to be sensitive to the point process realizations. A second open question is (Q2) can we design continuous and more robust non-parametric estimations, which account for the variance and the noise in the observed real life data?

The Bayesian inference for the Hawkes process has also been studied, including the work of Rasmussen [2013]; Linderman and Adams [2014]; Linderman and Adams [2015]. These works require either constructing a parametric triggering kernel [Rasmussen, 2013; Linderman and Adams, 2014] or discretizing the input domain to scale with the data size [Linderman and Adams, 2015]. To overcome the aforementioned shortcomings of discretization, Donnet et al. [2018] propose a continuous non-parametric Bayesian Hawkes process and resort to an unscalable Markov chain Monte Carlo (MCMC) approach to the posterior distribution. The third open question is (Q3) can we design efficient inference for non-parametric Bayesian Hawkes processes?

### 1.2 Thesis Contributions

In the thesis, we answer the above three questions by proposing four solutions:
(i) the Laplace Bayesian Hawkes process [Zhang et al., 2019], in Section 1.2.1,
(ii) the variational Bayesian Hawkes process [Zhang et al., 2020b], in Section 1.2.2,
(iii) the quantile propagation [Zhang et al., 2020a], in Section 1.2.3,
(iv) the kernel maximum moment restriction estimation [Zhang et al., 2021a], in Section 1.2.4.

The first two works, namely, (i) the Laplace Bayesian Hawkes process and (ii) the variational Bayesian Hawkes process, are proposed as direct solutions to all three questions (Q1~3). They are applications of existing approximation methods for Bayesian inference. Compared with the former, the latter employs the more complicated and advanced variational inference method, and enjoys faster training speed.

The last two works (iii) and (iv) study robust estimation methods beyond the field of the Hawkes process and are more general solutions to Q2~3. The work of quantile propagation explores a new approximation method for Bayesian inference. This method is similar to the expectation propagation (EP) algorithm [Opper and Winther, 2000] in the use of iterative local updates but employs the $L_{2}$ Wasserstein distance instead of the Kullback-Leibler (KL) divergence. Due to the locality of the computations, it is simple for our method and EP to parallelize and distribute, leading to more efficiency than variational inference [Li et al., 2015] especially on large-scale data [Gelman et al., 2017]. However unlike variational inference which minimizes the lower bound of the log model evidence, EP and QP correspond to no explicit global objective functions being minimized, and there is lack of works providing a clear understanding of the iterative updates, making EP and our QP behave more in a more complicated way than variational inference. Consequently, applying the new method to the Bayesian Hawkes process is challenging and non-trivial, and left for future work. It is observed that our method outperforms EP on the Gaussian process binary classification tasks, so we expect it to offer efficient and accurate performance for the Bayesian Hawkes process on large scale data.

Different from the above Bayesian solutions (i~iii), the final one (iv) proposes a frequentist estimation method which has an advantage of robustness and simplicity.

This method is a generalization of maximum likelihood estimation and the method of moments that are commonly employed by existing frequentist estimation methods [Lewis and Mohler, 2011; Zhou et al., 2013; Bacry and Muzy, 2016; Eichler et al., 2017]. Different from obtaining robustness via modeling uncertainty in the Bayesian methods, the method has the spirit of adversarial learning. It considers a reproducing kernel Hilbert space (RKHS) weighted average of the loss function values on data points and estimates the model parameters by optimizing the maximal or the worst average value. This idea is general and applicable to a wide range of estimation tasks, including estimation of the Hawkes process (as elaborated in Section 2.4), and we demonstrate its effectiveness on instrumental variable regression, which we are particularly interested in.

### 1.2.1 Laplace Bayesian Hawkes Process

In the first solution, we exploit block Gibbs sampling [Ishwaran and James, 2001] to iteratively sample the latent branching structure, the background intensity $\mu$ and the triggering kernel $\phi$. In each iteration, the point data are decomposed as a cluster of Poisson processes based on the sampled branching structure. The posterior $\mu$ and $\phi$ are estimated using the resulting cluster processes. Our framework is close to the stochastic Expectation-Maximization (EM) algorithm [Celeux and Diebolt, 1985] where posterior $\mu$ and $\phi$ are estimated [Lloyd et al., 2015; Walder and Bishop, 2017] in the M -step and random samples of $\mu$ and $\phi$ are drawn. We adapt the approach of the recent non-parametric Bayesian estimation for Poisson process intensities, termed Laplace Bayesian Poisson process [Walder and Bishop, 2017], to estimate the posterior $\phi$ given the sampled branching structure. We utilize the finite support assumption of the Hawkes Process to speed up sampling and to compute the probability of the branching structure. We theoretically show our method to be of linear time complexity. Furthermore, we explore the connection with the EM algorithm [Dempster et al., 1977] and develop a second variant of our method, as an approximate EM algorithm. We empirically show that our method enjoys linear time complexity and can infer known analytical kernels, i.e., exponential and sinusoidal kernels. On two large-scale social media datasets, our method outperforms the current state-of-the-art non-parametric approaches and the learned kernels reflect the perceived longevity for different content types. We propose a new acceleration trick based on the finite support assumption of the triggering kernel. The new trick enjoys higher efficiency than previous methods and accelerates the variational inference schema to linear time complexity per iteration.

### 1.2.2 Variational Bayesian Hawkes Process

The second solution proposes the first sparse Gaussian process modulated Hawkes process which employs a novel variational inference scheme, enjoys a linear time complexity per iteration and scales to large real world data. Our method is inspired by the variational Bayesian Poisson process (VBPP) [Lloyd et al., 2015] which provides Bayesian non-parametric inference only for the whole intensity of the Hawkes process
rather than its components: the background intensity $\mu$ and the triggering kernel $\phi$. Thus, the VBPP loses the internal reactions between points, and developing the variational Bayesian non-parametric inference for the Hawkes process is non-trivial and more challenging than the VBPP. In this work, we adapt the VBPP for the Hawkes process and term the new approach the variational Bayesian Hawkes process (VBHP). We employ a sparse Gaussian process modulated triggering kernel and a Gamma distributed background intensity and propose a new variational inference scheme for such models. Specifically, we employ the branching structure of the HP so that maximization of the evidence lower bound (ELBO) is tractable by the expectationmaximization algorithm, and we contribute a tighter ELBO which improves the fitting performance of our model. We propose a new acceleration trick based on the finite support assumption of the triggering kernel. The new trick enjoys higher efficiency than prior methods and accelerates the variational inference schema to linear time complexity per iteration. We empirically show that VBHP provides more accurate predictions than state-of-the-art methods on synthetic data and on two large online diffusion datasets. We validate the linear time complexity and faster convergence of our accelerated variational inference scheme compared to the Gibbs sampling method, and the practical utility of our tighter ELBO for model selection, which outperforms the common one in model selection.

### 1.2.3 Quantile Propagation for Gaussian Process Models

Beyond the area of the Hawkes process, we develop an approximate inference algorithm for Gaussian process models with factorized likelihoods based on minimization of the $L_{2}$ Wasserstein distance. Our approach employs a Gaussian likelihood to approximate the non-Gaussian likelihood. In order to optimize Gaussian likelihoods, we avoid directly minimizing the global $L^{2}$ Wasserstein distance between true and approximate joint posteriors, due to its computational and analytical intractability in high dimensional spaces. Instead, our method iteratively minimizes local $L^{2}$ Wasserstein distances, like EP. As a result, our method matches two quantile functions different from moment matching in EP, thus named quantile propagation. We further derive updating formulas for the mean and the variance of the Gaussian likelihood. The estimation of the mean is equal to EP's, while the variance is less than EP's, hence alleviating EP's deficiency of over-estimating variances [Minka, 2005; Heess et al., 2013; Hernández-Lobato et al., 2016]. We show that the optimal approximate Gaussian likelihood enjoys an economical parameterization like EP, i.e., relying on a single latent variable instead of all of them. This property allows our method or EP to significantly reduce memory consumption by a factor $N$ (the number of data) and computation time via optimizing much less $\left(O(1)\right.$, vs $O\left(N^{2}\right)$ for the full parameterization) parameters in each local update. We regard both methods as approximate coordinate descent algorithms to a KL divergence and a $L_{2}$ Wasserstein objective function respectively, under the same approximation assumption. In the experiment part, we evaluate EP and our method via Gaussian process binary classification on a number of real world datasets. Our results show that our method can outperform

EP in both predictive accuracy and uncertainty quantification, which validates our method alleviating over-estimation of the variance. We will apply this method to the Bayesian Hawkes process in the future and consider parallelizing or distributing the local updates. We expect it to perform efficiently and accurately on large-scale data.

### 1.2.4 Kernel Maximum Moment Restriction Estimation

The last solution proposes a simple framework for the nonlinear instrumental variable regression, which is a more difficult and general problem than the model parameter estimation for the Hawkes process. The framework is based on a kernelized conditional moment restriction known as a kernel maximum moment restriction (KMMR). The KMMR is formulated by maximizing the interaction between the residual and the instruments belonging to a unit ball in an RKHS. The KMMR allows us to reformulate the IV regression as a single-step empirical risk minimization problem, where the risk depends on the reproducing kernel on the instrument and can be estimated by a U-statistic or V-statistic. This simplification not only eases the proofs of consistency and asymptotic normality in both parametric and non-parametric settings, but also results in easy-to-use algorithms with an efficient hyper-parameter selection procedure. We demonstrate the advantages of our framework over existing methods using experiments on both synthetic and real-world data.

### 1.3 Broader Impact

The proposed inference techniques for Hawkes processes in the thesis have an advantage of a linear time complexity. Most of existing applications of Hawkes processes, from earthquake forecasting, finance to social media, take quadratic time complexity to estimate model parameters. Our algorithms thus will help to develop efficient applications. The approximate inference techniques also have the potential to be applied to other models to improve their robustness and account for uncertainty.

### 1.4 Thesis Outline

The subsequent chapters provide details of solutions to the proposed research questions. We first introduce prerequisites of the solutions and review related developments of them in Chapter 2. Then we elaborate on the four proposed methods in Chapters 3, 4, 5 and 6 respectively, which are summarized concisely in Chapter 7. Chapter 7 also provides some interesting future research directions.

## Chapter 2

## Preliminaries and Related Work

The works in the thesis are on Hawkes processes, approximate inference for Gaussian processes and robust frequentist estimation. In this chapter, we introduce the preliminaries of them and review the related works in the three research areas separately. Specifically, in Section 2.1, we first introduce the Poisson and Hawkes processes. We then review the Gaussian process model and the approximate Bayesian inference approaches used in the thesis in Section 2.2. In Section 2.3 and 2.4, we introduce the Wasserstein distance and the conditional moment restriction respectively. Then, we review the related work on estimation of Hawkes process, on expectation propagation and Wasserstein distance, and on condition moment restriction based estimation, in Sections 2.5, 2.6 and 2.7 respectively.

### 2.1 Poisson and Hawkes Processes

### 2.1.1 Poisson Processes

The Poisson (point) process assumes that points in any bounded subregion of the domain are independent of those in other subregions. More specifically, for any subset of the $d$-dimensional real space $\mathcal{T} \in \mathbb{R}^{d}$, other intervals $\mathcal{T}^{\prime}$ disjoint from $\mathcal{T}$ do not affect the occurrence of events in $\mathcal{T}$, and the probability of "the event count in $\mathcal{T}$, $N(\mathcal{T})$, being equal to $n^{\prime \prime}$ is determined by the Poisson distribution:

$$
P\{N(\mathcal{T})=n\}:=\frac{\Lambda(\mathcal{T})^{n}}{n!} e^{-\Lambda(\mathcal{T})}, \quad \Lambda(\mathcal{T}):=\int_{\mathcal{T}} \lambda(x) \mathrm{d} x
$$

where $\Lambda(\mathcal{T})$ is the expected number of points in $\mathcal{T}$, and the function $\lambda(x)$ is known as the intensity and modulates the occurrence rate of events at $x$. The log-likelihood of $\mathcal{D}:=\left\{x_{i}\right\}_{i=1}^{N}$ given $\lambda$ is [Rubin, 1972]:

$$
\log p(\mathcal{D} \mid \lambda)=\sum_{i=1}^{N} \log \lambda\left(x_{i}\right)-\int_{\Omega} \lambda(x) \mathrm{d} x,
$$

where $\Omega$ is the sample domain of $\left\{x_{i}\right\}_{i=1}^{N}$. The log-likelihood of any stochastic point process, such as the Hawkes process, has the same form. We recommend [Daley and Vere-Jones, 2003] for a detailed introduction to stochastic point processes.


Figure 2.1: A Poisson process realization (points). The intensity function $\lambda(t)=10 \sin (t)+$ 10.5 (solid line) is used to simulate points and higher function values generates more points.

### 2.1.2 Hawkes Processes

The Hawkes process [Hawkes, 1971] is a self-exciting point process, in which the occurrence of a point increases the arrival rate $\lambda(\cdot)$ of new points. Given a set of ordered points $\mathcal{D}=\left\{x_{i}\right\}_{i=1}^{N}, x_{i} \in \mathbb{R}^{d}$, the intensity at $x$ conditioned on given points is written as:

$$
\lambda(x)=\mu+\sum_{x_{i}<x} \phi\left(x-x_{i}\right),
$$

where $\mu>0$ is background intensity, commonly considered as a constant, and $\phi: \mathbb{R}^{d} \rightarrow[0, \infty)$ is the triggering kernel. We consider $d=1$ for a concise presentation and extension to $d>1$ follows by the same development procedure. In the thesis, we are interested in the branching structure of the Hawkes process. As introduced in Chapter 1, each point $x_{i}$ has a parent that we represent by the below one-hot vector. The index is up to $i-1$ as the points are ordered.

$$
\begin{equation*}
\boldsymbol{b}_{i}=\left[b_{i 0}, b_{i 1}, \cdots, b_{i, i-1}\right]^{T} . \tag{2.1}
\end{equation*}
$$

Each element $b_{i j}$ is binary, $b_{i j}=1$ represents that $x_{i}$ is triggered by $x_{j}(0 \leq j \leq i-1$, $x_{0}$ : the background), and $\sum_{j=0}^{i-1} b_{i j}=1$. A branching structure $B$ is the set of $\boldsymbol{b}_{i}$ and is determined by $B=\left\{\boldsymbol{b}_{i}\right\}_{i=1}^{N}$. We define the probability of $b_{i j}=1$ as (see e.g. [Lewis and Mohler, 2011])

$$
p_{i j}:=p\left(b_{i j}=1\right)= \begin{cases}\phi\left(x_{i}-x_{j}\right) / \lambda\left(x_{i}\right), & 0<j \leq j-1  \tag{2.2}\\ \mu / \lambda\left(x_{i}\right), & j=0\end{cases}
$$

then the probability of $B$ is expressed as below and it is clear $\sum_{j=0}^{i-1} p_{i j}=1$ for all $i$.

$$
p(B)=\prod_{i=1}^{N} \prod_{j=0}^{i-1} p_{i j}^{b_{i j}}
$$

Log-likelihood Given Branching Structure. Given the branching structure $B$, the Hawkes process can be viewed as a cluster of Poisson processes, namely, $\operatorname{PP}(\mu)$ and $\left\{\operatorname{PP}\left(\phi\left(x-x_{i}\right)\right)\right\}_{i=1}^{N}$ as introduced in Chapter 1. Consequently, the log likelihood of the Hawkes process becomes a sum of log-likelihoods of Poisson processes. The data domain of $\operatorname{PP}(\mu)$ equals that of the Hawkes process, which we denote as $\Omega$, and we denote the data domain of $\operatorname{PP}\left(\phi\left(x-x_{i}\right)\right)$ by $\Omega_{i} \subset \Omega$. As a result, the log-likelihood of data points and the branching structure is calculated by:

$$
\begin{equation*}
\log p(\mathcal{D}, B \mid \mu, \phi)=\sum_{i=1}^{N}\left(\sum_{j=1}^{i-1} b_{i j} \log \phi_{i j}+b_{i 0} \log \mu\right)-\sum_{i=1}^{N} \int_{\Omega_{i}} \phi-\mu|\mathcal{T}|, \tag{2.3}
\end{equation*}
$$

where $|\Omega|:=\int_{\Omega} 1 \mathrm{~d} x, \phi_{i j}:=\phi\left(x_{i}-x_{j}\right)$ and $\int_{\Omega_{i}} \phi:=\int_{\Omega_{i}} \phi(x) \mathrm{d} x$. Note that the Lebesgue measure is considered here.


Figure 2.2: A Hawkes process realization (points). The intensity function $\lambda(t)=1+$ $\sum_{t_{i}<t} 0.1 e^{-\left(t-t_{i}\right)}+0.6 e^{-3\left(t-t_{i}\right)}+0.7 e^{-7\left(t-t_{i}\right)}$ (solid line) is used to simulate points.

### 2.2 Gaussian Processes

### 2.2.1 Exact Gaussian Processes

The Gaussian process (GP) is defined as a distribution over a function $f$, denoted as $\mathcal{G} \mathcal{P}(f)$, such that for any non-empty set of $x:=\left\{x_{1}, \cdots, x_{N}\right\}$ in the domain of $f$, function values $f:=\left\{f\left(x_{i}\right)\right\}_{i=1}^{N}$ jointly have a Gaussian distribution, i.e.,

$$
\begin{equation*}
p(\boldsymbol{f} \mid \theta)=\mathcal{N}\left(\boldsymbol{\mu}_{x}, \boldsymbol{K}_{x}\right), \tag{2.4}
\end{equation*}
$$

where the mean vector $\mu_{x}$ and the covariance matrix $K_{x}$ are the evaluation of a mean function $\mu(\cdot)$ and a covariance function $k(\cdot, \cdot)$ on $\boldsymbol{x}$, that is, $\boldsymbol{\mu}_{x}:=\left\{\mu\left(x_{i}\right)\right\}_{i=1}^{N}$ and $\boldsymbol{K}_{x}:=\left[k\left(x_{i}, x_{j}\right)\right]_{i, j=1}^{N}$, and $\theta$ is the set of parameters of $\mu$ and $k$. The covariance function defines the covariance between two function values,

$$
k(x, y)=\operatorname{Cov}(f(x), f(y)),
$$

and any covariance function $k$ can be represented in terms of the eigenvalues $\left\{\lambda_{i}\right\}_{i=1}^{K}$ and eigenfunctions $\left\{e_{i}\right\}_{i=1}^{K}$ according to Mercer's theorem [Mercer, 1909]

$$
k(x, y)=\sum_{i=1}^{K} \lambda_{i} e_{i}(x) e_{i}(y)
$$

where $\left\{e_{i}\right\}_{i=1}^{K}$ are chosen to be orthonormal in $L^{2}(\Omega, v)$ for some sample space $\Omega$ with the measure $v$ and $K=\infty$ if $k$ is non-degenerate. A popular approach to determine $\theta$ is to maximize the $\log$ marginal likelihood given a set of observations $\mathcal{D}$,

$$
\begin{equation*}
\widehat{\theta}=\underset{\theta}{\operatorname{argmax}} \log p(\mathcal{D} \mid \theta):=\log \int p(\mathcal{D} \mid \boldsymbol{f}, \theta) p(\boldsymbol{f} \mid \theta) \mathrm{d} f . \tag{2.5}
\end{equation*}
$$

With the optimized parameters $\widehat{\theta}$, the posterior distribution of $f$ is obtained as

$$
\begin{equation*}
p(\boldsymbol{f} \mid \mathcal{D}, \widehat{\theta})=\frac{p(\mathcal{D} \mid \boldsymbol{f}, \widehat{\theta}) p(\boldsymbol{f} \mid \widehat{\theta})}{\int p(\mathcal{D} \mid \boldsymbol{f}, \widehat{\theta}) p(\boldsymbol{f} \mid \widehat{\theta}) \mathrm{d} \boldsymbol{f}} \tag{2.6}
\end{equation*}
$$

This exact GP model often suffers from computational and analytical intractability. Specifically, computing the exact posterior requires to store and invert an $(N \times N)$ matrix, which consumes $O\left(N^{2}\right)$ memory and $O\left(N^{3}\right)$ time. Both impede the method from scaling up to large problems. The analytical intractability originates from the non-conjugate Gaussian prior $p(\boldsymbol{f} \mid \theta)$ for the likelihood $p(\mathcal{D} \mid f, \theta)$ and as a results, the integral in Equation (2.6) (as well as Equation (2.5)) has no closed-form expression. To overcome the two issues, the sparse GP model (Section 2.2.2) and approximate Bayesian inference (Section 2.2.3, 2.2.4, 2.2.5) are often employed.

### 2.2.2 Sparse Gaussian Processes

The sparse GP [Quiñonero-Candela and Rasmussen, 2005; Titsias, 2009] is proposed to reduced the computational burden. With the same assumption on function distributions as that of the GP, the sparse GP introduces inducing points to approximate distributions of function values at any point, which is realized based on the Bayes' rules for Gaussian variables. As a result, the computational time and memory are reduced to $O\left(N M^{2}\right)$ and $O(N M)$ respectively, where $M$ is the number of inducing points.

More specifically, given a set of $M$ inducing points $z:=\left\{z_{1}, \cdots, z_{M}\right\} \subset \Omega$ and corresponding function values $u:=\left\{f\left(z_{i}\right)\right\}_{i=1}^{M}, f$ evaluated at $x$ have a joint Gaussian distribution

$$
\binom{f}{u} \sim \mathcal{N}\left(\binom{\boldsymbol{\mu}_{x}}{\boldsymbol{\mu}_{z}},\left(\begin{array}{cc}
\boldsymbol{K}_{x} & \boldsymbol{K}_{x z} \\
\boldsymbol{K}_{z x} & \boldsymbol{K}_{z}
\end{array}\right)\right)
$$

where $K_{x}$ and $K_{z}$ are the covariance matrices of $x$ and $z$ respectively, as defined in Equation (2.4), and $K_{x z}$ and $K_{z x}$ are cross-covariance matrices between $x$ and $z$, which are obtained as the evaluations of the covariance function on $x$ and $z$, namely,
$\boldsymbol{K}_{x z}=\left\{k\left(x_{i}, z_{j}\right)\right\}_{i, j=1}^{N, M}$ and $\boldsymbol{K}_{x z}^{\top}=\boldsymbol{K}_{z x}$. The conditional distribution of $\boldsymbol{f}$ given $\boldsymbol{u}$ can then be expressed as

$$
p(\boldsymbol{f} \mid \boldsymbol{u})=\mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{x}}+\boldsymbol{K}_{x z} \boldsymbol{K}_{z}^{-1}\left(\boldsymbol{u}-\boldsymbol{\mu}_{z}\right), \boldsymbol{K}_{\boldsymbol{x}}-\boldsymbol{K}_{x z} \boldsymbol{K}_{z}^{-1} \boldsymbol{K}_{z x}\right) .
$$

We omit the condition on $\theta$ from time to time for a neat presentation. The quantity of interest is the posterior distribution of $f$, which can be calculated based on the posterior distribution of $u$ as

$$
p(\boldsymbol{f} \mid \mathcal{D})=\int p(\boldsymbol{f} \mid \boldsymbol{u}) p(\boldsymbol{u} \mid \mathcal{D}) \mathrm{d} \boldsymbol{u}
$$

where $\mathcal{D}$ is the observation. The posterior distribution of $u$ often has no closed-form expressions and the Sparse GP model employs a Gaussian distribution $q(\boldsymbol{u})$ for $\boldsymbol{u}$, which is an approximation to the posterior distribution of $\boldsymbol{u}$. Suppose that $q(\boldsymbol{u})$ has a Gaussian form

$$
q(\boldsymbol{u})=\mathcal{N}(\boldsymbol{m}, \boldsymbol{S}),
$$

and then the optimal $m$ and $S$ of $q(\boldsymbol{u})$ can be obtained by using variational inference for the marginal likelihood (Equation (2.5)) [Titsias, 2009]. With the optimal $q(\boldsymbol{u}) \approx p(\boldsymbol{u} \mid \mathcal{D})$, we obtain a distribution over $f$, which is Gaussian again and is the approximation of interest to the posterior distribution of $f$,

$$
\begin{aligned}
q(\boldsymbol{f}) & =\int p(\boldsymbol{f} \mid \boldsymbol{u}) q(\boldsymbol{u}) \mathrm{d} \boldsymbol{u} \\
& =\mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{x}}+\boldsymbol{K}_{x z} \boldsymbol{K}_{z}^{-1}\left(\boldsymbol{m}-\boldsymbol{\mu}_{z}\right), \boldsymbol{K}_{x}-\boldsymbol{K}_{x z} \boldsymbol{K}_{z}^{-1}\left(\boldsymbol{K}_{z}-\boldsymbol{S}\right) \boldsymbol{K}_{z}^{-1} \boldsymbol{K}_{z x}\right) \\
& \approx \int p(\boldsymbol{f} \mid \boldsymbol{u}) p(\boldsymbol{u} \mid \mathcal{D}) \mathrm{d} \boldsymbol{u} \\
& =p(\boldsymbol{f} \mid \mathcal{D}) .
\end{aligned}
$$

### 2.2.3 Laplace Approximation

Laplace approximation [Rasmussen and Williams, 2005, section 3.4], [Bishop, 2006, section 4.4] is a widely-used method and it provides a Gaussian approximation to the analytical intractable posterior probability. Suppose that the analytically intractable probability has the following form,

$$
p(\boldsymbol{f})=Z^{-1} h(\boldsymbol{f}), \quad Z:=\int h(\boldsymbol{f}) \mathrm{d} f
$$

where $Z$ is the unknown normalization constant. First, the method finds a mode of $p(f)$, which is a point $f_{0}$ satisfying

$$
\left.\nabla h(f)\right|_{f=f_{0}}=\mathbf{0} .
$$

Then, it considers a second-order Taylor expansion of $\ln h(f)$ centred on the mode $f_{0}$,

$$
\ln h(f) \approx \ln h\left(f_{0}\right)-\frac{1}{2}\left(f-f_{0}\right)^{\top} A\left(f-f_{0}\right), \quad A:=\left.\nabla^{2} \ln h(f)\right|_{f=f_{0}}
$$

where the first-order term is omitted as it is equal to zero. Taking the exponential and normalization on both sides of the above equation, we obtain a Gaussian distribution approximating $p(f)$,

$$
p(f) \approx \frac{|A|^{1 / 2}}{(2 \pi)^{N / 2}} \exp \left\{-\frac{1}{2}\left(f-f_{0}\right)^{\top} A\left(f-f_{0}\right)\right\}=\mathcal{N}\left(f_{0}, A^{-1}\right)
$$

### 2.2.4 Variational Inference

Variational inference [Bishop, 2006, Section 10.1] is a more general approach than Laplace approximation because it allows non-Gaussian approximation to intractable distributions. Consider the joint distribution of observations and variables, $p(\mathcal{D}, f \mid \theta)$. The variational approach introduces a variational distribution $q\left(f \mid \theta^{\prime}\right)$ to approximate the posterior distribution $p(f \mid \mathcal{D}, \theta)$ and optimizes $q\left(f \mid \theta^{\prime}\right)$ by maximizing a lower bound of the log-likelihood, known as the evidence lower bound (ELBO), which can be derived from the non-negative gap perspective:

$$
\begin{align*}
\log p(\mathcal{D} \mid \theta) & =\log \frac{p(\mathcal{D}, f \mid \theta)}{q\left(f \mid \theta^{\prime}\right)}-\log \frac{p(f \mid \mathcal{D}, \theta)}{q\left(f \mid \theta^{\prime}\right)} \\
& =\underbrace{\underbrace{\mathbb{E}_{q\left(f \mid \theta^{\prime}\right)}[\log p(\mathcal{D} \mid f, \theta)]}_{\text {reconstruction term }}-\underbrace{\operatorname{KL}\left(q\left(f \mid \theta^{\prime}\right)| | p(f \mid \theta)\right)}_{\text {regularization term }}}_{\equiv \operatorname{ELBO}(q(f), p(\mathcal{D} \mid f), p(f))}+\underbrace{\operatorname{KL}\left(q\left(f \mid \theta^{\prime}\right) \| p(f \mid \mathcal{D}, \theta)\right)}_{\text {intractable (non-negative) gap }} \\
& \geq \operatorname{ELBO}(q(f), p(\mathcal{D} \mid \boldsymbol{f}), p(f)), \tag{2.7}
\end{align*}
$$

where we omit $\theta$ and $\theta^{\prime}$ in conditions. For notational convenience, we will often omit conditioning on $\theta$ and $\theta^{\prime}$ hereinafter. Optimizing the ELBO w.r.t. $\theta^{\prime}$ balances between the reconstruction error and the Kullback-Leibler (KL) divergence from the prior. Generally, the conditional $p(\mathcal{D} \mid \boldsymbol{f})$ is known, so is the prior. Thus, for an appropriate choice of $q$, it is easier to work with this lower bound than with the intractable posterior $p(f \mid \mathcal{D})$. We also see that, due to the form of the intractable gap, if $q$ is from a distribution family containing elements close to the true unknown posterior, then $q$ will be close to the true posterior when the ELBO is close to the true likelihood. An alternative derivation applies Jensen's inequality [Jordan et al., 1999]. A vanilla ELBO objective can usually be straight-forwardly minimized by standard algorithms, such as gradient descent. In Chapter 4, we will show that the ELBO of our Bayesian Hawkes process model has extra constraints on certain variables, so we accordingly develop a variational-inference based two-step iterative optimization algorithm.

### 2.2.5 Expectational Propagation

In this subsection, we introduce the application of the expectational propagation (EP) algorithm to the GP models [Opper and Winther, 2000; Minka, 2001b,c]. Given $\mathcal{D}=\left\{y_{i}\right\}_{i=1}^{N}$ and $f_{i}:=f\left(x_{i}\right)$, suppose that the GP model has the factorized likelihood,

$$
p(\mathcal{D} \mid \boldsymbol{f})=\prod_{i=1}^{N} p\left(y_{i} \mid f_{i}\right)
$$

Numerous problems take this form: binary classification [Williams and Barber, 1998], single-output regression with Gaussian likelihood [Matheron, 1963], Student's-t likelihood [Jylänki et al., 2011] or Poisson likelihood [Zou, 2004], and the warped GP [Snelson et al., 2004]. EP deals with the analytical intractability by using Gaussian approximations to the individual non-Gaussian likelihoods, namely,

$$
p\left(y_{i} \mid f_{i}\right) \approx t_{i}\left(f_{i}\right) \equiv \widetilde{Z}_{i} \mathcal{N}\left(f_{i} \mid \widetilde{\mu}_{i}, \widetilde{\sigma}_{i}^{2}\right) .
$$

The function $t_{i}$ is often called the site function and is specified by the site parameters: the scale $\widetilde{Z}_{i}$, the mean $\widetilde{\mu}_{i}$ and the variance $\widetilde{\sigma}_{i}^{2}$. Notably, it is sufficient to use univariate site functions given that the local update can be efficiently computed using the marginal distribution only [Seeger, 2005]. We refer to this as the locality property. Although in this thesis we employ a more complex $L_{2}$ Wasserstein distance, our approach retains this property, as we elaborate in Chapter 5.

Given the site functions, one can approximate the intractable posterior distribution $p(f \mid \mathcal{D})$ using a Gaussian $q(f)$ as below,

$$
\begin{align*}
q(\boldsymbol{f} \mid \mathcal{D}) & =q(\mathcal{D})^{-1} p(\boldsymbol{f}) \prod_{i=1}^{N} t_{i}\left(f_{i}\right) \equiv \mathcal{N}(\boldsymbol{f} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})  \tag{2.8}\\
\boldsymbol{\mu} & =\boldsymbol{\Sigma}\left(\boldsymbol{K}_{x}^{-1} \boldsymbol{\mu}_{x}+\widetilde{\boldsymbol{\Sigma}}^{-1} \widetilde{\boldsymbol{\mu}}\right), \quad \boldsymbol{\Sigma}=\left(\boldsymbol{K}_{x}^{-1}+\widetilde{\boldsymbol{\Sigma}}^{-1}\right)^{-1}
\end{align*}
$$

where $\widetilde{\mu}$ is the vector of $\widetilde{\mu}_{i}, \widetilde{\boldsymbol{\Sigma}}$ is diagonal with $\widetilde{\boldsymbol{\Sigma}}_{i i}=\widetilde{\sigma}_{i}^{2} ; \log q(\mathcal{D})$ is the log approximate model evidence expressed as below and further employed to optimize the GP hyperparameters:

$$
\begin{equation*}
\log q(\mathcal{D})=\sum_{i=1}^{N} \log \left(\widetilde{Z}_{i} / \sqrt{2 \pi}\right)-\frac{1}{2} \log \left|\boldsymbol{K}_{x}+\widetilde{\boldsymbol{\Sigma}}\right|-\frac{1}{2} \widetilde{\boldsymbol{\mu}}^{\mathrm{T}}\left(\boldsymbol{K}_{x}+\widetilde{\boldsymbol{\Sigma}}\right)^{-1} \widetilde{\boldsymbol{\mu}} . \tag{2.9}
\end{equation*}
$$

The core of EP is to optimize site functions $\left\{t_{i}\left(f_{i}\right)\right\}_{i=1}^{N}$. Ideally, one would seek to minimize the (global) KL divergence between the true and approximate posterior distributions $\operatorname{KL}(p(f \mid \mathcal{D}) \| q(f))$, however this is intractable. Instead, EP is built based on the assumption that the global divergence can be approximated by the local divergence $\operatorname{KL}(\widetilde{q}(f) \| q(f))$, where $\widetilde{q}(f) \propto q^{\backslash i}(f) p\left(y_{i} \mid f_{i}\right)$ and $q^{\backslash i}(f) \propto q(f) / t_{i}\left(f_{i}\right)$ are referred to as the tilted and cavity distributions, respectively. Note that the cavity distribution is Gaussian while the tilted distribution is usually not. The local divergence can be simplified from multi-dimensional to univariate, $\operatorname{KL}(\widetilde{q}(f) \| q(f))=\operatorname{KL}\left(\widetilde{q}\left(f_{i}\right) \| q\left(f_{i}\right)\right)$
(detailed in Appendix C.7), and $t_{i}\left(f_{i}\right)$ is optimized by minimizing it.
The minimization is realized by projecting the tilted distribution $\widetilde{q}\left(f_{i}\right)$ onto the Gaussian family, with the projected Gaussian denoted as

$$
\operatorname{proj}_{\mathrm{KL}}\left(\widetilde{q}\left(f_{i}\right)\right):=\underset{\mathcal{N}}{\operatorname{argmin}} \mathrm{KL}\left(\widetilde{q}\left(f_{i}\right) \| \mathcal{N}\left(f_{i}\right)\right) .
$$

Then the projected Gaussian is used to update $t_{i}\left(f_{i}\right) \propto \operatorname{proj}_{\mathrm{KL}}\left(\widetilde{q}\left(f_{i}\right)\right) / q^{\backslash i}\left(f_{i}\right)$. The mean and the variance of $\operatorname{proj}_{\mathrm{KL}}\left(\widetilde{q}\left(f_{i}\right)\right) \equiv \mathcal{N}\left(\mu^{\star}, \sigma^{\star 2}\right)$ match the moments of $\widetilde{q}\left(f_{i}\right)$ and are used to update $t_{i}\left(f_{i}\right)$ 's parameters:

$$
\begin{align*}
& \mu^{\star}=\mu_{\widetilde{q}_{i}} \quad \sigma^{\star 2}=\sigma_{\tilde{q}_{i^{\prime}}}^{2}  \tag{2.10}\\
& \widetilde{\mu}_{i}=\widetilde{\sigma}_{i}^{2}\left(\mu^{\star}\left(\sigma^{\star}\right)^{-2}-\mu_{\backslash i} \sigma_{\backslash i}^{-2}\right), \widetilde{\sigma}_{i}^{-2}=\left(\sigma^{\star}\right)^{-2}-\sigma_{\backslash i}^{-2}, \tag{2.11}
\end{align*}
$$

where $\mu_{\widetilde{q}_{i}}$ and $\sigma_{\widetilde{q}_{i}}^{2}$ are the mean and the variance of $\widetilde{q}\left(f_{i}\right)$, and $\mu_{\backslash i}$ and $\sigma_{\backslash i}^{2}$ are the mean and the variance of $q^{\backslash i}\left(f_{i}\right)$. We refer to the projection as the local update. Note that $\widetilde{Z}$ does not impact the optimization of $q(f)$ or the GP hyper-parameters $\theta$, so we omit the update formula for $\widetilde{Z}$. We summarize EP in Algorithm 2 (Appendix).

### 2.3 Wasserstein Distance

We denote by $\mathcal{M}_{+}^{1}(\Omega)$ the set of all probability measures on $\Omega$. We consider probability measures on the $d$-dimensional real space $\mathbb{R}^{d}$. The Wasserstein distance between two probability distributions $\xi, v \in \mathcal{M}_{+}^{1}\left(\mathbb{R}^{d}\right)$ can be intuitively defined as the cost of transporting the probability mass from one distribution to the other. We are particularly interested in the subclass of $L_{p}$ Wasserstein distance, formally defined as follows.

Definition 1 ( $L_{p}$ Wasserstein distance). Consider the set of all probability measures on the product space $\mathbb{R}^{d} \times \mathbb{R}^{d}$, whose marginal measures are $\xi$ and $v$ respectively, denoted as $U(\xi, v)$. The $L_{p}$ Wasserstein distance between $\xi$ and $v$ is defined as

$$
W_{p}^{p}(\xi, v):=\inf _{\pi \in U(\xi, v)} \int_{\mathbb{R}^{d} \times \mathbb{R}^{d}}\|x-z\|_{p}^{p} d \pi(x, z),
$$

where $p \in[1, \infty)$ and $\|\cdot\|_{p}$ is the $L_{p}$ norm.
Like the KL divergence, the $L_{p}$ Wasserstein distance it has a minimum of zero, achieved when the distributions are equivalent. Unlike the $K L$, however, it is a proper distance metric, and thereby satisfies the triangle inequality, and has the appealing property of symmetry.

A less fundamental property of the Wasserstein distance which we exploit for computational efficiency is:

Proposition 1. [Peyré et al., 2019, Remark 2.30] The $L_{p}$ Wasserstein distance between 1-d distribution functions $\xi$ and $v \in \mathcal{M}_{+}^{1}(\mathbb{R})$ equals the $L_{p}$ distance between the quantile
functions,

$$
W_{p}^{p}(\xi, v)=\int_{0}^{1}\left|F_{\tilde{\zeta}}^{-1}(y)-F_{v}^{-1}(y)\right|^{p} d y
$$

where $F_{z}: \mathbb{R} \rightarrow[0,1]$ is the cumulative distribution function (CDF) of $z$, defined as $F_{z}(x)=$ $\int_{-\infty}^{x} d z$, and $F_{z}^{-1}$ is the pseudo-inverse or quantile function, defined as $F_{z}^{-1}(y)=\min _{x}\{x \in$ $\left.\mathbb{R} \cup\{-\infty\}: F_{z}(x) \geq y\right\}$.

Finally, the following translation property of the $L_{2}$ Wasserstein distance is central to our proof of locality in Chapter 5:

Proposition 2. [Peyré et al., 2019, Remark 2.19] Consider the $L_{2}$ Wasserstein distance defined for $\xi$ and $v \in \mathcal{M}_{+}^{1}\left(\mathbb{R}^{d}\right)$, and let $f_{\tau}(\boldsymbol{x})=\boldsymbol{x}-\boldsymbol{\tau}, \boldsymbol{\tau} \in \mathbb{R}^{d}$, be a translation operator. If $\xi_{\tau}$ and $v_{\boldsymbol{\tau}^{\prime}}$ denote the probability measures of translated random variables $f_{\tau}(\boldsymbol{x}), \boldsymbol{x} \sim \xi$, and $f_{\boldsymbol{\tau}^{\prime}}(\boldsymbol{x})$, $x \sim v$, respectively, then

$$
W_{2}^{2}\left(\xi_{\tau}, v_{\boldsymbol{\tau}^{\prime}}\right)=W_{2}^{2}(\xi, v)-2\left(\boldsymbol{\tau}-\boldsymbol{\tau}^{\prime}\right)^{T}\left(\boldsymbol{m}_{\xi}-\boldsymbol{m}_{v}\right)+\left\|\boldsymbol{\tau}-\boldsymbol{\tau}^{\prime}\right\|_{2}^{2},
$$

where $\boldsymbol{m}_{\tilde{\xi}}$ and $\boldsymbol{m}_{v}$ are means of $\xi$ and $v$ respectively. In particular when $\boldsymbol{\tau}=\boldsymbol{m}_{\tilde{\xi}}$ and $\boldsymbol{\tau}^{\prime}=\boldsymbol{m}_{v}$, $\xi_{\tau}$ and $v_{\tau^{\prime}}$ become zero-mean measures, and

$$
W_{2}^{2}\left(\xi \tau, v_{\tau^{\prime}}\right)=W_{2}^{2}(\xi, v)-\left\|\boldsymbol{m}_{\xi}-\boldsymbol{m}_{v}\right\|_{2}^{2} .
$$

### 2.4 Conditional Moment Restriction

Let $(X, Z)$ be a random variable taking values in $\mathcal{X} \times \mathcal{Z}$ and $\Theta$ a parameter space. A conditional moment restriction (CMR) [Newey, 1993; Ai and Chen, 2003] can then be expressed as

$$
\begin{equation*}
\operatorname{CMR}\left(\theta_{0}\right)=\mathbb{E}\left[\varphi_{\theta_{0}}(X) \mid Z\right]=\mathbf{0}, \quad P_{Z}-\text { almost surely (a.s.) } \tag{2.12}
\end{equation*}
$$

for the true parameter $\theta_{0} \in \Theta$. The function $\varphi_{\theta}(X)$ is a problem-dependent generalized residual function in $\mathbb{R}^{q}$ parameterized by $\theta$. Intuitively, the CMR asserts that, for correctly specified models, the conditional mean of the generalized residual function is almost surely equal to zero. Many statistical models can be written as Equation (2.12) including nonparametric regression models where $X=(\tilde{X}, Y), Z=\tilde{X}$ and $\varphi_{\theta}(X)=Y-f(\tilde{X} ; \theta)$; conditional quantile models where $X=(\tilde{X}, Y), Z=\tilde{X}$, and $\varphi_{\theta}(X)=\mathbb{1}\{Y<f(\tilde{X} ; \theta)\}-\tau$ for the target quantile $\tau \in[0,1] ;$ IV regression models where $X=(\tilde{X}, Y), Z$ is an IV, and $\varphi_{\theta}(X)=Y-f(\tilde{X} ; \theta)$; the stochastic point process, where $X=(\tilde{X}), Z$ is the observed region in the domain, and $\varphi_{\theta}(X)=\nabla_{\theta} \log p_{\theta}(X)$ with $p_{\theta}$ the probability model of the point process. More specifically in the point process case, the maximum likelihood estimator is written in general as below

$$
\hat{\theta}=\left.\underset{\theta}{\operatorname{argmax}} \mathbb{E}_{X \sim p}\left[\log p_{\theta}(X) \mid Z\right] \Longrightarrow \mathbb{E}_{X \sim p}\left[\nabla_{\theta} \log p_{\theta}(X) \mid Z\right]\right|_{\theta=\hat{\theta}}=\mathbf{0},
$$



Figure 2.3: A causal graph depicting an instrumental variable $Z$ that satisfies an exclusion restriction and unconfoundedness (there may be a confounder $\varepsilon$ acting on $X$ and $Y$, but it is independent of $Z$ ).
where $p$ is the true probability of data, and as we note, $\hat{\theta}$ is often estimated as the solution to the first-order condition shown in the right part. Besides, the finite region Z where real-world data are collected is random and dependent on the time when we observe the data, so it is reasonable to view the point process as a CMR model.

### 2.4.1 Instrumental Variable Regression

We introduce details of instrumental variable (IV) regression on which we test our kernel maximum moment restriction estimation method. Following standard setting in the literature [Hartford et al., 2017; Lewis and Syrgkanis, 2018; Bennett et al., 2019; Singh et al., 2019; Muandet et al., 2020b], let $X$ be a treatment (endogenous) variable taking value in $\mathcal{X} \subseteq \mathbb{R}^{d}$ and $Y$ a real-valued outcome variable. The goal of IV regression is to estimate a function $f: \mathcal{X} \rightarrow \mathbb{R}$ from a structural equation model (SEM) of the form

$$
\begin{equation*}
Y=f(X)+\varepsilon, \quad X=t(Z)+g(\varepsilon)+v \tag{2.13}
\end{equation*}
$$

where we assume that $\mathbb{E}[\varepsilon]=0$ and $\mathbb{E}[v]=0$. Unfortunately, as we can see from Equation (2.13), $\varepsilon$ is correlated with the treatment $X$, i.e., $\mathbb{E}[\varepsilon \mid X] \neq 0$, and hence standard regression methods cannot be used to estimate $f$. This setting arises, for example, when there exist unobserved confounders (i.e., common causes) between $X$ and $Y$.

To illustrate the problem, let us consider an example taken from Hartford et al. [2017] which aims at predicting sales of airline ticket $Y$ under an intervention in price of the ticket $X$. However, there exist unobserved variables that may affect both sales and ticket price, e.g., conferences, COVID-19 pandemic, etc. This creates a correlation between $\varepsilon$ and $X$ in (2.13) that prevents us from applying the standard regression toolboxes directly on observational data.

Instrumental variable (IV). To address this problem, we assume access to an instrumental variable $Z$ taking value in $\mathcal{Z} \subseteq \mathbb{R}^{d^{\prime}}$. As we can see in (2.13), the instrument $Z$ is associated with the treatments $X$, but not with the outcome $Y$, other than through its effect on the treatments. Formally, $Z$ must satisfy (i) Relevance: $Z$ has a causal influence on $X$, (ii) Exclusion restriction: $Z$ affects $Y$ only through $X$, i.e., $Y \Perp Z \mid X, \varepsilon$, and (iii) Unconfounded instrument(s): $Z$ is independent of the error, i.e., $\varepsilon \Perp Z$. For example, the instrument $Z$ may be the cost of fuel, which influences sales only via price. Intuitively, $Z$ acts as a "natural experiment" that induces variation in $X$; see Figure 2.3.

### 2.5 Related Works of Hawkes Process Estimation

As mentioned in Chapter 1, the triggering kernel $\phi$ is important as it is shared and decides the class of the whole process. In this section, we review the estimation of the triggering kernel $\phi$ as well as topics related to it, including modeling of the intensity function and of the integral of the triggering/intensity function, and model parameter estimation methods and so on.

### 2.5.1 Parametric Frequentist Solutions

The Hawkes processes with the parametric triggering kernels are applied in a wide range of areas and the popularity can be explained by their simplicity and flexibility compared with other models. Mishra et al. [2016] employ the branching factor of the Hawkes process with the power-law kernel to predict popularity of tweets; Kurashima et al. [2018] predict human actions using a Hawkes process equipped with exponential, Weibull and Gaussian mixture kernels; online popularity unpredictability is explained using the Hawkes process with a variant of the exponential kernel by Rizoiu et al. [2018]; Xu et al. [2016] employ a sum of Gaussian kernels to discover the Granger causality for the Hawkes process. The model parameters can be easily determined by the maximum likelihood estimation [Ozaki, 1979]. However, most works employing Hawkes processes with parametric triggering kernels encode strong assumptions, and limit the expressivity of the models. Therefore, recent works design practical approaches to learn flexible representations of the optimal triggering kernel from data, as the following subsections.

### 2.5.2 Non-parametric Frequentist Solutions

A popular direction on learning the flexible triggering kernel functions is the nonparametric frequentist estimation. Basically, the triggering kernel function is assigned a non-parametric form which is mainly a piecewise function [Lewis and Mohler, 2011; Zhou et al., 2013; Bacry and Muzy, 2016; Eichler et al., 2017]. The form of the function is caused by discretization of the function domain and leads to poor scaling with the domain dimension and sensitivity to the choice of discretization. Moreover, they do not quantify the uncertainty of the learned triggering kernels, so are sensitive to randomness of finite data. In contrast, our methods require no discretization and are Bayesian, so has an advantage of scalability with the dimension of the domain and robustness to uncertainty in the finite samples.

Frequentist Estimation of Intensities. Learning the triggering kernels based on the maximum likelihood requires $O\left(N^{2}\right)$ computational time and to circumvent this problem, many frequentist methods focus on estimating the intensity function, which needs $O(N)$ time. These methods rely on modern machine learning models, such as the neural network [Du et al., 2016; Xiao et al., 2017a,b; Mei and Eisner, 2017; Omi et al., 2019; Zhang et al., 2020; Zuo et al., 2020] and the reproducing kernel Hilbert space [Flaxman et al., 2017] (RKHS). As no assumption is made on the type of interaction between points such as self-excitement, they allow the modelling of flexible point
processes. Early neural network approaches on modeling intensity functions rely on the recurrent neural network (RNN) or long short-term memory (LSTM) network [Du et al., 2016; Xiao et al., 2017b; Mei and Eisner, 2017], and estimate model parameters based on maximum likelihood. Training deep RNNs and LSTMs is notoriously difficult because of gradient explosion and gradient vanishing [Pascanu et al., 2013]. The inherently sequential nature of RNN and LSTM renders the impossibility to process all the events in parallel and limits the methods' ability to scale to large datasets. More recent works deal with this issue by employing the transformer or the self-attention mechanism [Zhang et al., 2020; Zuo et al., 2020]. Apart from the issues with models, most of the neural network approaches suffer from no analytical expression for the integral of the intensity function in the likelihood. Hence, the integral has to be approximated numerically, which leads to in-accurate parameter estimation. In order to circumvent this problem, Omi et al. [2019] directly model the integral by a RNN. As a result, the intensity function is obtained by differentiation of the integral, which is easier to compute than integration. Shchur et al. [2020b] present a new modeling method for the integral by exploiting the normalizing flows which later is employed to model the distribution of the inter-arrival time [Shchur et al., 2020a]. As shown by Walder and Bishop [2017], the integral can have a closed-form expression with kernel methods by explicitly constructing the eigen-components of the kernel function. Orthogonal to the afore-mentioned development, there is recent interest in parameter estimation methods beyond the popular maximum likelihood estimation, including the Wiener-hopf equation based method [Bacry and Muzy, 2016], the adversarial loss such as the wasserstein distance based loss [Xiao et al., 2017a], the least square loss [Eichler et al., 2017], the cumulants-based method [Achab et al., 2017], the reinforcement learning based method [Li et al., 2018]. Interestingly, it is unnecessary to estimate the triggering kernel or the intensity function in some applications, such as discovering the Granger causality, where it is sufficient to only estimate the integral of the triggering kernel [Achab et al., 2017]. We recommend the related work sections of these methods for more development on point processes. Different from works reviewed here, this thesis focuses on the non-parametric Bayesian estimation of the triggering kernel.

### 2.5.3 Bayesian Parametric and Non-parametric Solutions

The Bayesian non-parametric estimation for the Hawkes process has been studied, including the work of Rasmussen [2013]; Linderman and Adams [2014]; Linderman and Adams [2015]. These work require either constructing a parametric triggering kernel [Rasmussen, 2013; Linderman and Adams, 2014] or discretizing the input domain to scale with the data size [Linderman and Adams, 2015]. The shortcoming of discretization is just mentioned and to overcome it, Donnet et al. [2018] propose a continuous non-parametric Bayesian Hawkes process and resort to an unscalable Markov chain Monte Carlo (MCMC) estimator to the posterior distribution. We comparatively summarize a part of related works in 2.1. Compared with these works, our methods are Bayesian and non-parametric without requiring discretization of

Table 2.1: Non-parametric triggering kernel estimation.
$\left.\begin{array}{lllll}\hline \text { Methods } & \begin{array}{l}\text { Time } \\ \text { Complexity }\end{array} & & \text { Bayesian } & \text { Continuous }\end{array} \begin{array}{l}\text { Non- } \\ \text { parametric }\end{array}\right]$
domains and has an advantage of a linear time complexity allowing it to be applied to large-scale real-world datasets.

### 2.6 Related Works of Expectation Propagation and Wasserstein Distance

The basis of the EP algorithm for GP models was first proposed by Opper and Winther [2000] and then generalized by Minka [2001b,c]. Power EP [Minka, 2004, 2005] is an extension of EP that exploits the more general $\alpha$-divergence (with $\alpha=1$ corresponding to the forward KL divergence in EP) and has been recently used in conjunction with GP pseudo-input approximations [Bui et al., 2017]. Although generally not guaranteed to converge locally or globally, Power EP uses fixed-point iterations for its local updates and has been shown to perform well in practice for GP regression and classification [Bui et al., 2017]. In comparison, our approach (quantile propagation) uses the $L_{2}$ Wasserstein distance, and like EP, it yields convex local optimizations for GP models with factorized likelihoods. This convexity benefits the convergence of the local update, and is retained even with the general $L_{p}(p \geq 1)$ Wasserstein distance as shown in Theorem 2 (Chapter 5). Moreover, for the same class of GP models, both EP and our approach have the locality property [Seeger, 2005] and can be unified in the generic message passing framework [Minka, 2005].

Without the guarantee of convergence for the explicit global objective function, understanding EP has proven to be a challenging task. As a result, a number of works have instead attempted to directly minimize divergences between the true and approximate joint posteriors, for divergences such as the KL [Jordan et al., 1999; Dezfouli and Bonilla, 2015], Rényi [Li and Turner, 2016], $\alpha$ [Hernández-Lobato et al., 2016] and optimal transport divergences [Ambrogioni et al., 2018]. To deal with the infinity issue of the KL (and more generally the Rényi and $\alpha$ divergences) which arises from different distribution supports [Montavon et al., 2016; Arjovsky et al., 2017; Gulrajani et al., 2017], Hensman et al. [2014] employ the product of tilted distributions
as an approximation. A number of variants of EP have also been proposed, including the convergent double loop algorithm [Opper and Winther, 2005], parallel EP [Minka, 2001a], distributed EP built on partitioned datasets [Xu et al., 2014; Gelman et al., 2017], averaged EP assuming that all approximate likelihoods contribute similarly [Dehaene and Barthelmé, 2018], and stochastic EP which may be regarded as sequential averaged EP [Li et al., 2015].

The $L_{2}$ Wasserstein distance between two Gaussian distributions has a closed form expression [Dowson and Landau, 1982]. Detailed research on the Wasserstein geometry of the Gaussian distribution is conducted by Takatsu [2011]. Recently, this closed form expression has been applied to robust Kalman filtering [ShafieezadehAbadeh et al., 2018] and to the analysis of populations of GPs [Mallasto and Feragen, 2017]. A more general extension to elliptically contoured distributions is provided by Gelbrich [1990] and used to compute probabilistic word embeddings [Muzellec and Cuturi, 2018]. A geometric interpretation for the $L_{2}$ Wasserstein distance between any distributions [Benamou and Brenier, 2000] has already been exploited to develop approximate Bayesian inference schemes [El Moselhy and Marzouk, 2012]. Our approach is based on the $L_{2}$ Wasserstein distance but does not exploit these closed form expressions; instead we obtain computational efficiency by leveraging the EP framework and using the quantile function form of the $L_{2}$ Wasserstein distance for univariate distributions. We believe our work paves the way for further practical approaches to Wasserstein-distance-based Bayesian inference.

### 2.7 Related Works of Conditional Moment Restriction

We review conditional moment restriction (CMR) models [Newey, 1993; Ai and Chen, 2003; Dikkala et al., 2020] in this section and the models have a wide range of applications in causal inference, economics, and finance modeling, where for correctly-specified models the conditional mean of certain functions of data equals zero almost surely. This kind of models also appear in Mendelian randomization, a technique in genetic epidemiology that uses genetic variation to improve causal inference of a modifiable exposure on disease [Davey Smith and Ebrahim, 2003; Burgess et al., 2017a]. Rational expectation models [Muth, 1961], widely-used in macroeconomics, measures how available information is exploited to form future expectations by decision-makers as conditional moments [Muth, 1961]. Furthermore, CMRs have also gained popularity in the community of causal machine learning, leading to novel algorithms such as generalized random forests [Athey et al., 2019], double/debiased machine learning [Chernozhukov et al., 2018] and nonparametric IV regression [Bennett et al., 2019; Muandet et al., 2020b]; see also related works therein, as well as in offline reinforcement learning [Liao et al., 2021].

This thesis focuses on CMR based estimation with the application to IV regression (introduced in Section 2.4). Therefore, we mainly review literature on the CMR based estimation for IV regression.

Classical methods for IV regression often rely on a linearity assumption in which
a two-stage least squares (2SLS) is the most popular technique [Angrist et al., 1996]. The generalized method of moments (GMM) of Hansen [1982], which imposes the orthogonality restrictions, can also be used for the linear IV regression. For nonlinear regression, numerous methodologies have been developed in the field of nonparametric IV [Newey and Powell, 2003; Hall and Horowitz, 2005; Blundell et al., 2007; Horowitz, 2011] and recent machine learning [Hartford et al., 2017; Lewis and Syrgkanis, 2018; Bennett et al., 2019; Muandet et al., 2020b; Singh et al., 2019]. However, these estimators have complicated structures in general. The nonparametric IV is an ill-posed problem or requires estimating a conditional density. The machine learning methods needs two-stage estimation or minimax optimization. As a result, it is not easy to obtain an asymptotic distribution of estimation errors and to simply apply learning algorithms such as the stochastic gradient descent (SGD), e.g., see Daskalakis and Panageas [2018] for the limitation of SGD with minimax problems.

Several extensions of 2SLS and GMM exist for the nonlinear IV problem. In the two-stage approach, the function $f(x)$ has often been obtained by solving a Fredholm integral equation of the first kind $\mathbb{E}[Y \mid Z]=\int f(x) \mathrm{d} P(X \mid Z)$. In Newey and Powell [2003]; Blundell et al. [2007]; Horowitz [2011]; Chen and Pouzo [2012], linear regression is replaced by a linear projection onto a set of known basis functions. A uniform convergence rate of this approach is provided in Chen and Christensen [2018]. However, it remains an open question how to best choose the set of basis functions. In Hall and Horowitz [2005] and Darolles et al. [2011], the first-stage regression is replaced by a conditional density estimation of $\mathbb{P}(X \mid Z)$ using a kernel density estimator. Estimating conditional densities is a difficult task and is known to perform poorly in high dimension [Tsybakov, 2008].

The IV regression has also recently received attention in the machine learning community. Hartford et al. [2017] proposed to solve the integral equation by first estimating $P(X \mid Z)$ with a mixture of deep generative models on which the function $f(x)$ can be learned with another deep NNs. Instead of NNs, Singh et al. [2019] proposed to model the first-stage regression using the conditional mean embedding of $P(X \mid Z)$ [Song et al., 2009, 2013; Muandet et al., 2017] which is then used in the second-step kernel ridge regression. In other words, the first-stage estimation in Singh et al. [2019] becomes a vector-valued regression problem. The critical drawback of these algorithms is that they involve the intermediate first-stage regression which may not be of our primary interest. In an attempt to alleviate this drawback, Muandet et al. [2020b] and Liao et al. [2020] reformulate the two-stage procedure as a convex-concave saddle-point problem. The DualIV [Muandet et al., 2020b] has a quadratic objective function similar to ours, but its RKHS is applied on $(Z, Y)$ which cannot be interpreted as a valid instrument. In contrast, the approach of Liao et al. [2020, Appendix F] is highly related to GMM and provides a dual reformulation of our method by using a RKHS for the inner maximization. Besides, the starting points of these works differ from ours. In both DualIV and Liao et al. [2020], they started from the population risk functional, whereas we start from the CMR. The fact that they arrive at the same objective highlights a deeper connection which requires further investigation.

Our work follows in spirit many GMM-based approaches for IV regression, namely,

Lewis and Syrgkanis [2018]; Bennett et al. [2019]; Muandet et al. [2020a]. We apply the MMR framework of Muandet et al. [2020a] to the IV regression problem, whereas Muandet et al. [2020a] only considers a conditional moment (CM) testing problem. In fact, this framework was initially inspired by Lewis and Syrgkanis [2018] and Bennett et al. [2019] which instead parametrize the instruments by deep NNs and Muandet et al. [2020b] which proposes the dual formulation of the two-stage procedure. By combining the GMM framework with RKHS functions, our objective function can be evaluated in closed-form. As a result, our IV estimate can be obtained by minimizing the empirical risk, as opposed to an adversarial optimization used in Lewis and Syrgkanis [2018] and Bennett et al. [2019]. Furthermore, unlike existing two-stage procedures [Angrist et al., 1996; Hartford et al., 2017; Singh et al., 2019], our algorithm does not require the first-stage regression. It is important to note that, concurrent to our work, Dikkala et al. [2020] extends the work of Lewis and Syrgkanis [2018] to an algorithm similar to our MMR-IV (RKHS) [Dikkala et al., 2020, Section 4]. Although both work employs RKHSs in the minimax frameworks, Dikkala et al. [2020] incorporate a Tikhonov regularization on $h$ in (6.2) and resort to the representer theorem [Schölkopf et al., 2001a] to develop the analytical objective function, whereas we impose a unit-ball constraint which is a form of Ivanov regularization [Ivanov et al., 2002] and enables not to rely on such a theorem.

Beyond the IV regression, there are numerous prior studies that are related to ours, especially in policy evaluation, reinforcement learning, and causal inference. We leave the review of them to Appendix D.4.

### 2.8 Summary

In this chapter, we first review different technical prerequisites, including the Poisson and Hawkes processes, the Gaussian process model, the approximate Bayesian inference approaches, the Wasserstein distance and the conditional moment restriction respectively. We then review the related development of these techniques. The prerequisite are essential for the methods proposed in the following four chapters. In the next chapter, we present an efficient non-parametric Bayesian inference framework for estimation of the Hawkes process. The approach is developed based on the branching structure of the Hawkes process and exploit Gibbs sampling and Laplace approximation.

# Gibbs Sampling and Laplace Approximation Based Efficient Inference 

In this chapter, we present a general framework for the efficient non-parametric Bayesian inference of Hawkes processes. The contents are categorized as:
(i) In Section 3.1, we first review a prerequisite of our work, the Laplace Bayesian Poisson process.
(ii) In Section 3.2, we propose an efficient non-parametric Bayesian framework for the Hawkes process by combining Gibbs sampling and the Laplace Bayesian Poisson process. We exploit block Gibbs sampling [Ishwaran and James, 2001] to iteratively sample the latent branching structure, the background intensity $\mu$ and the triggering kernel $\phi$ from their posterior distributions. In each iteration, the point data are decomposed as a cluster of Poisson processes based on the sampled branching structure and the posterior distributions of $\mu$ and $\phi$ are estimated using the resulting cluster processes. Our framework is close to the stochastic Expectation-Maximization (EM) algorithm [Celeux and Diebolt, 1985] where posterior $\mu$ and $\phi$ are estimated [Lloyd et al., 2015; Walder and Bishop, 2017] in the M-step and random samples of $\mu$ and $\phi$ are drawn. We adapt the approach of the recent non-parametric Bayesian estimation for Poisson process intensities, termed Laplace Bayesian Poisson process (LBPP) [Walder and Bishop, 2017], to estimate the posterior $\phi$ given the sampled branching structure. Especially in Section 3.2.4, we utilize the finite support assumption of the Hawkes Process to speed up sampling and computing the probability of the branching structure. We theoretically show our method to be of linear time complexity.
(iii) In Section 3.3, we furthermore explore the connection with the EM algorithm [Dempster et al., 1977] and develop a second variant of our method, as an approximate EM algorithm.
(iv) In Section 3.4, we empirically show our method enjoys linear time complexity and can infer known analytical kernels, i.e., exponential and sinusoidal kernels. On two large-scale social media datasets, our method outperforms the current
state-of-the-art non-parametric approaches and the learned kernels reflect the perceived longevity for different content types.

### 3.1 Laplace Bayesian Poisson Process

In this section, we introduce a prerequisite of our work, the Laplace Bayesian Poisson process (LBPP) [Walder and Bishop, 2017]. LBPP has been proposed for the nonparametric Bayesian estimation of the intensity of a Poisson process. To satisfy non-negativity of the intensity function, LBPP models the intensity function $\lambda$ as a permanental process [Shirai and Takahashi, 2003], i.e., $\lambda=g \circ f$ where the link function $g(z)=z^{2} / 2$ and $f(\cdot)$ obeys a Gaussian process (GP) prior. Alternative link functions include $\exp (\cdot)$ [Møller et al., 1998; Diggle et al., 2013] and $g(z)=$ $\lambda^{*}(1+\exp (-z))^{-1}$ [Adams et al., 2009] where $\lambda^{*}$ is constant.

The choice $g(z)=z^{2} / 2$ has the analytical advantages; for some covariances the log-likelihood can be computed in closed form [Lloyd et al., 2015; Flaxman et al., 2017]. LBPP exploits the Mercer expansion [Mercer, 1909] of the GP covariance function $k(x, y) \equiv \operatorname{Cov}(f(x), f(y))$, namely,

$$
\begin{equation*}
k(x, y)=\sum_{i=1}^{K} \lambda_{i} e_{i}(x) e_{i}(y) \tag{3.1}
\end{equation*}
$$

where for non-degenerate kernels, $K=\infty$. The eigenfunctions $\left\{e_{i}(\cdot)\right\}_{i=1}^{K}$ are chosen to be orthonormal in $L^{2}(\Omega, \mathcal{M})$ for some sample space $\Omega$ with the measure $\mathcal{M} . f(\cdot)$ can be represented as a linear combination of $e_{i}(\cdot)$ [Rasmussen and Williams, 2005, section 2.2] as below

$$
\begin{equation*}
f(\cdot)=\omega^{T} e(\cdot), \quad \omega \sim \mathcal{N}(0, \Lambda) \tag{3.2}
\end{equation*}
$$

where $\Lambda=\operatorname{diag}\left(\lambda_{1}, \lambda_{2}, \cdots, \lambda_{K}\right)$ is a diagonal covariance matrix and $\boldsymbol{e}(\cdot)=\left\{e_{1}(\cdot)\right\}_{i=1}^{K}$ is a (column) vector of basis functions. Computing the posterior distribution of the intensity function $\lambda$ is equivalent to estimating the posterior distribution of $\omega$ which, in LBPP, is approximated by a normal distribution (as known as Laplace approximation introduced in Section 2.2.3). That is

$$
\log p(\omega \mid x, \Omega, k) \approx \log \mathcal{N}(\omega \mid \hat{\omega}, Q)
$$

where $x:=\left\{x_{i}\right\}_{i=1}^{N}$ is a set of point data, $\Omega$ the sample space and $k$ the Gaussian process kernel function. $\hat{\omega}$ is selected as the mode of the true posterior and Q the negative inverse Hessian of the true posterior at $\hat{w}$ :

$$
\begin{align*}
\hat{\boldsymbol{\omega}} & =\underset{\omega}{\operatorname{argmax}} \log p(\boldsymbol{\omega} \mid x, \Omega, k),  \tag{3.3}\\
Q^{-1} & =-\left.\partial_{\omega \omega^{T}} \log p(\omega \mid x, \Omega, k)\right|_{\omega=\hat{\omega}} . \tag{3.4}
\end{align*}
$$

The approximate posterior distribution of $f(x)$ is expressed as the following normal
distribution [Rasmussen and Williams, 2005, section 2.2]:

$$
\begin{equation*}
f(x) \sim \mathcal{N}\left(\hat{\boldsymbol{\omega}}^{T} \boldsymbol{e}(x), \boldsymbol{e}(x)^{T} Q \boldsymbol{e}(x)\right) \equiv \mathcal{N}\left(\nu, \sigma^{2}\right) . \tag{3.5}
\end{equation*}
$$

Furthermore, the posterior distribution of $\lambda(x)=f(x)^{2} / 2$ is a Gamma distribution:

$$
\begin{equation*}
\operatorname{Gamma}(t \mid \alpha, \beta):=\beta^{\alpha} t^{\alpha-1} e^{-\beta t} / \Gamma(\alpha), \tag{3.6}
\end{equation*}
$$

where $\alpha=\left(v^{2}+\sigma^{2}\right)^{2} /\left(4 v^{2} \sigma^{2}+2 \sigma^{4}\right)$ and $\beta=\left(v^{2}+\sigma^{2}\right) /\left(2 v^{2} \sigma^{2}+\sigma^{4}\right)$.

### 3.2 Inference via Sampling

We now detail our efficient non-parametric Bayesian estimation algorithm for Hawkes processes, which employs block Gibbs sampling to iteratively draw samples from the posterior distributions of $\mu$ (constant background intensity) and $\phi(\cdot)$ (triggering kernel). Our method starts with random $\mu_{0}$ and $\phi_{0}(\cdot)$, and iterates by cycling through the following four steps ( $k$ is the iteration index):
(i) Calculate $p\left(B \mid x, \phi_{k-1}, \mu_{k-1}\right)$, the distribution of the branching structure $B$ given the data $\boldsymbol{x}$, triggering kernel $\phi_{k-1}$, and background intensity $\mu_{k-1}$ (see details in Section 3.2.1).
(ii) Sample a branching structure $B_{k}$ from $p\left(B \mid x, \phi_{k-1}, \mu_{k-1}\right)$ (Section 3.2.1).
(iii) Estimate $p\left(\phi \mid B_{k}, \boldsymbol{x}\right)$ (Section 3.2.3) and $p\left(\mu \mid B_{k}, \boldsymbol{x}\right)$ (Section 3.2.2).
(iv) Sample a sample $\phi_{k}$ and $\mu_{k}$ from $p\left(\phi \mid B_{k}, x\right)$ and $p\left(\mu \mid B_{k}, \boldsymbol{x}\right)$, respectively.

By standard Gibbs sampling arguments, the samples of $\phi$ and $\mu$ drawn in the step (iv) converge to the desired posterior, modulo the Laplace approximation for estimation of $p\left(\phi \mid B_{k}, x\right)$ in (iii). As the method is based on block Gibbs sampling [Ishwaran and James, 2001], we term it Gibbs-Hawkes in this chapter.

### 3.2.1 Distribution and Sampling of the Branching Structure

The branching structure $B$ has a data structure of tree (as Figure 1.1(b)) and consists of independent triggering events. Therefore, we may sample a branching structure by sampling a parent for each $x_{i}$ independently, where sampling exploits the probabilities of triggering events for $x_{i}$, namely, $\left\{p_{i j}\right\}_{j=0}^{i-1}$ as defined in Equation (2.2). The sampled branching structure separates a set of points into immigrants and offspring (introduced in Chapter 1). Immigrants can be regarded as a sequence generated from $\operatorname{PP}(\mu)$, where $\operatorname{PP}(\mu)$ is a Poisson process owning an intensity $\mu$, and can be used to estimate the posterior distribution of $\mu$.

The key property which we exploit in the subsequent Section 3.2.2 and Section 3.2.3 is the following. Denote by $\left\{x_{k}^{(i)}\right\}_{k=1}^{N_{x_{i}}}$, the $N_{x_{i}}$ offspring is generated by point $x_{i}$. If such a sequence is aligned to an origin at $x_{i}$, yielding $S_{x_{i}}:=\left\{t_{k}^{(i)}-x_{i}\right\}_{k=1}^{N_{i}}$, then the aligned sequence is drawn from $\operatorname{PP}(\phi)$ over $\left[0, \mathrm{~T}-x_{i}\right]$ where $[0, T]$ is the sample domain of the Hawkes process. The posterior distribution of $\phi$ is estimated on all such aligned sequences.

### 3.2.2 Posterior Distribution of $\mu$

Continuing from the observations in Section 3.2.1, note that if we are given a set of points $\left\{x_{i}\right\}_{i=1}^{M}(M \leq N)$ generated by $\operatorname{PP}(\mu)$ over $\Omega=[0, T]$, the likelihood for $\left\{x_{i}\right\}_{i=1}^{M}$ is the Poisson likelihood, $p\left(\left\{x_{i}\right\}_{i=1}^{M} \mid \mu, \Omega\right)=e^{-\mu T}(\mu T)^{M} / M$ !. For simplicity, we place a conjugate (Gamma) prior on $\mu T, \mu T \sim \operatorname{Gamma}(\alpha, \beta)$; the Gamma-Poisson conjugate family conveniently gives the posterior distribution of $\mu T$, i.e., $p\left(\mu T \mid\left\{x_{i}\right\}_{i=1}^{M}, \alpha, \beta\right)=$ $\operatorname{Gamma}(\alpha+M, \beta+1)$. We choose the scale $\alpha$ and the rate $\beta$ in the Gamma prior by making the mean of the Gamma posterior equal to $N$ and the variance $M / 2$, which is easily shown to correspond to $\alpha=M$ and $\beta=1$. Finally, due to conjugacy we obtain the posterior

$$
p\left(\mu \mid\left\{x_{i}\right\}_{i=1}^{M}, \alpha, \beta\right)=\operatorname{Gamma}(2 M, 2 T) .
$$

### 3.2.3 Posterior Distribution of $\phi$

We handle the posterior distribution of the triggering kernel $\phi$ given the branching structure in an analogous manner to the LBPP method of Walder and Bishop [2017]. That is, we assume that $\phi(\cdot)=f^{2}(\cdot) / 2$ where $f(\cdot)$ is Gaussian process distributed as described in Section 3.1. In line with [Walder and Bishop, 2017], we consider the sample domain $[0, \pi]$ and the so-called cosine kernel,

$$
\begin{align*}
k(x, y) & =\sum_{\gamma \geq \mathbf{0}} \lambda_{\gamma} e_{\gamma}(x) e_{\gamma}(y),  \tag{3.7}\\
\lambda_{\gamma} & :=1 /\left(a\left(\gamma^{2}\right)^{m}+b\right),  \tag{3.8}\\
e_{\gamma}(x) & :=(2 / \pi)^{1 / 2} \sqrt{1 / 2}^{[\gamma=0]} \cos (\gamma x) . \tag{3.9}
\end{align*}
$$

Here, $\gamma$ is a multi-index with non-negative (integral) values, [•] is the indicator function, $a$ and $b$ are parameters controlling the prior smoothness, and we let $m=2$. The choice of $m$ affects the rate of change or shape of $\lambda_{\gamma}$ with $\gamma$ and results in different priors: for large $m \lambda_{\gamma}$ decreases rapidly with $\gamma$, giving a-priori preference to smoother functions. So do the parameters $a$ and $b$. This basis is orthonormal w.r.t. the Lebesgue measure on $\Omega=[0, \pi]$. The expansion Equation (3.7) is an explicit kernel construction based on the Mercer expansion as per Equation (3.1), but other kernels may be used, for example by Nyström approximation of the Mercer decomposition [Flaxman et al., 2017].

As mentioned at the end of Section 3.2.1, by conditioning on the branching structure we may estimate $\phi$ by considering the aligned sequences. In particular, letting $S_{x_{i}}$ denote the aligned sequence generated by $x_{i}$, the joint distribution of $\omega$ and $S:=\left\{S_{x_{i}}\right\}_{i=1}^{N}$ is calculated as [Walder and Bishop, 2017],

$$
\begin{align*}
& \log p(\omega, S \mid \Omega, k)=\sum_{i=1}^{N} \sum_{\Delta t \in S_{x_{i}}} \log \frac{1}{2}\left(\boldsymbol{\omega}^{T} \boldsymbol{e}(\Delta t)\right)^{2}-\frac{1}{2} \boldsymbol{\omega}^{T}\left(A+\Lambda^{-1}\right) \boldsymbol{\omega}+C,  \tag{3.10}\\
& A:=\sum_{i=1}^{N} \int_{0}^{T-x_{i}} \boldsymbol{e}(t) \boldsymbol{e}(t)^{T} \mathrm{~d} t, \quad C:=-\frac{1}{2} \log \left[(2 \pi)^{K}|\Lambda|\right],
\end{align*}
$$

where $K$ is the number of eigenfunctions and $\Lambda$ is defined in Equation (3.2). Note that there is a subtle but important difference between the integral term above and that of Walder and Bishop [2017], namely, the limit of integration; closed-form expressions for the present case are provided in Section A. 1 of the appendix. Putting the above equation into Equation (3.3) and Equation (3.4), and we obtain the mean $\hat{\boldsymbol{\omega}}$ and the covariance $Q$ of the (Laplace) approximate log-posterior in $\omega$ :

$$
\begin{align*}
\hat{\boldsymbol{\omega}} & =\underset{\omega}{\operatorname{argmax}} \log p(\boldsymbol{\omega}, \boldsymbol{S} \mid \Omega, k),  \tag{3.11}\\
Q^{-1} & =-\sum_{i=1}^{N} \sum_{\Delta t \in S_{x_{i}}} 2 \boldsymbol{e}(\Delta t) \boldsymbol{e}(\Delta t)^{T} /\left(\hat{\boldsymbol{\omega}}^{T} \boldsymbol{e}(\Delta t)\right)^{2}+A+\Lambda^{-1} . \tag{3.12}
\end{align*}
$$

Then, the posterior $\phi$ is achieved by Equation (3.5) and (3.6).

### 3.2.4 Computational Complexity

For the LBPP method, constructing Equation (3.10) and (3.12) takes $O\left(N_{o} K^{2}\right)$ where $K$ is the number of basis functions and $N_{o}$ is the number of offspring. Optimizing $\boldsymbol{\omega}$ (Equation (3.11)) is a concave problem, which can be solved efficiently. If LBFGS is used, $O(C K)$ will be taken to calculate the gradient on each $\omega$ where $C$ is the number of steps stored in memory. Computing $Q$ requires inverting a $K \times K$ matrix, which is $O\left(K^{3}\right)$. As a result, the complexity of estimating the conditional probability $p(\phi \mid B)$ is $O\left(\left(N_{o}+K\right) K^{2}\right)$. In terms of estimating $p(\mu \mid B)$ taking $O(1)$, the complexity of estimating $p(\mu \mid B)$ and $p(\phi \mid B)$ is linear to the number of data. The time taken to sample $\mu$ and $\phi$ is minor $(O(1)$ and $O(K)$ respectively), so estimation time dominates. Although the naive complexity for $p_{i j}$ is $O\left(N^{2}\right)$, Halpin [2012] provides an optimized approach to reduce it to $O(N)$, which relies on the finite support assumption of Hawkes processes. The finite support assumption says that the value of the triggering kernel is negligible when the input is large [Halpin, 2012, p. 9]. As a result, the step of sampling branching structures can also be run in $O(N)$ and points with negligible impacts on another point are not sampled as its parents. Interestingly, in comparison with LBPP, while our model is in some sense more complex, it enjoys a more favorable computational complexity. In summary, we have the following complexities per iteration and in Section 3.4, we validate the complexity on both synthetic and real data.

Table 3.1: Time complexity.

| Operation | $p(\mu \mid B)$ | $p_{i j}$ | $p(\phi \mid B)$ | overall |
| :---: | :---: | :---: | :---: | :---: |
| Complexity | $O(1)$ | $O(N)$ | $O\left(\left(N_{o}+K\right) K^{2}\right)$ | $O\left((N+K) K^{2}\right)$ |



Figure 3.1: A visual summary of the Gibbs-Hawkes, EM-Hawkes and the EM algorithms. The differences between them are (1) the number of sampled branching structures and (2) selected $\phi$ and $\mu$ for $p_{i j}$. In contrast with with Gibbs-Hawkes, the EM-Hawkes method draws multiple branching structures at once and calculates $p_{i j}$ using M.A.P. $\phi$ and $\mu$. The EM algorithm is equivalent to sampling infinite branching structures and exploiting M.A.P. or constrained M.L.E. $\phi$ and $\mu$ to calculate $p_{i j}$ (see Section 3.3).

### 3.3 Maximum-A-Posterior Estimation

We explore a connection between the sampler of Section 3.2 and the EM algorithm, which allows us to introduce an analogous but intermediate scheme between them. In contrast to the random sampler of Section 3.2, the proposed scheme employs a deterministic maximum-a-posteriori (M.A.P.) sampler.

### 3.3.1 Relationship to EM

At the very beginning of this chapter, we mentioned the connection between our method and the stochastic EM algorithm [Celeux and Diebolt, 1985]. The difference is in the M-step; to perform EM [Dempster et al., 1977], we need only modify our sampler by: (a) sampling infinite branching structures at each iteration, and (b) recalculating the probability of the branching structure with the M.A.P. $\mu$ and $\phi$, given the infinite set of branching structures. More specifically, maximizing the expected log posterior distribution to estimate M.A.P. $\mu$ and $\phi$ given infinite branching structures is equivalent to maximizing the EM objective in the M-step (see Section A. 2 of the appendix for the detailed derivation). Finally, note that the above step (b) is identical to the E-step of the EM algorithm.

### 3.3.2 EM-Hawkes

Following the discussion above, we propose EM-Hawkes, an approximate EM algorithm variant of Gibbs-Hawkes proposed in Section 3.2. Specifically, at each iteration EM-Hawkes (a) samples a finite number of cluster assignments (to approximate the


Figure 3.2: Computation time (seconds) for calculating $p_{i j}$ and sampling branching structures, with and without Halpin's speed up.
expected $\log$ posterior distribution), and (b) finds the M.A.P. triggering kernels and background intensities rather than sampling them as per block Gibbs sampling (the M-step of the EM algorithm). An overview of the Gibbs-Hawkes, EM-Hawkes and EM algorithm is illustrated in Figure 3.1.

Note that under our LBPP-like posterior, finding the most likely triggering kernel $\phi$ is intractable (see details in Appendix A.3). As an approximation, we take the element-wise mode of the marginal distributions of $\left\{\phi\left(x_{i}\right)\right\}_{i=1}^{N}$ to approximate the mode of the joint distribution of the $\left\{\phi\left(x_{i}\right)\right\}_{i=1}^{N}$.

### 3.4 Experiments

We now evaluate our proposed approaches - Gibbs-Hawkes and EM-Hawkes - and compare them to three baseline models, on synthetic data and on two large Twitter online diffusion datasets. The three baselines are:
(i) A naive parametric Hawkes equipped with a constant background intensity and an exponential (Exp) triggering kernel $\phi=a_{1} a_{2} \exp \left(-a_{2} t\right), a_{1}, a_{2}>0$, estimated by maximum likelihood.
(ii) Ordinary differential equation (ODE)-based non-parametric non-bayesian Hawkes [Zhou et al., 2013].
(iii) Wiener-Hopf (WH) equation based non-parametric non-bayesian Hawkes [Bacry and Muzy, 2016]. Codes of ODE based and WH based methods are publicly available [Bacry et al., 2017].

### 3.4.1 Synthetic Data

We employ two toy Hawkes processes to generate data, both having the same background intensity $\mu=10$, and cosine (Equation (3.13)) and exponential (Equation (3.14)) triggering kernels respectively. Note that compared to the cosine triggering kernel, the exponential one has a larger L2 norm for its derivative, and the difference is designed to test the performance of the approaches in different situations. We can check that both triggering kernels have negligible values when the input is large in the domain which we will choose as $[0, \pi]$, so the finite support assumption is satisfied.

$$
\begin{align*}
& \phi_{\cos }(x)=\cos (3 \pi x)+1, \quad x \in[0,1] ; \quad 0, \quad \text { otherwise; }  \tag{3.13}\\
& \phi_{\exp }(x)=5 \exp (-5 x), \quad x \geq 0 \tag{3.14}
\end{align*}
$$

Prediction. For three baseline models and EM-Hawkes, the predictions $\mu_{\text {pred }}$ and $\phi_{\text {pred }}$ are taken to be the M.A.P. values, while for Gibbs-Hawkes we use the posterior mean.

Evaluation. Each toy model generates 400 point sequences over $\Omega=[0, \pi]$, which are evenly split into 40 groups, 20 for training and 20 for test. Each of the three methods fit on each group, i.e., summing log-likelihoods for 10 sequences (for the parametric Hawkes) or estimating the log posterior probability of the Hawkes process given 10 sequences (for Gibbs-Hawkes and EM-Hawkes) or fitting the superposition of 10 sequences [ Xu et al., 2018]. Since the true models are known, we evaluate fitting results using the relative L2 distance between predicted and true $\mu$ and $\phi(\cdot)$ : $d_{L 2}\left(g_{\text {pred }}, g_{\text {true }}\right)=\left(\int_{\Omega}\left(g_{\text {pred }}(t)-g_{\text {true }}(t)\right)^{2} d t\right)^{1 / 2} /\left(\int_{\Omega}\left(g_{\text {true }}(t)\right)^{2} d t\right)^{1 / 2}$.

Table 3.2: Empirical performance comparison between algorithms (columns) with different measures (rows). Top: relative L2 distance to known $\phi$ and $\mu$, and AVG denoted the average of L2 errors of $\phi$ and $\mu$. bottom: mean predictive log likelihood on real data. Bold numbers denote the best performance and the underlined numbers for the second best.

| Data | Exp | ODE | WH | Gibbs | EM |
| :--- | :--- | :--- | :--- | :--- | :--- |
| $\phi_{\text {cos }}$ | 0.661 | 0.553 | 1.000 | 0.338 | 0.318 |
| $\mu_{\text {cos }}$ | 0.069 | 0.071 | 1.739 | 0.078 | 0.119 |
| AVG $_{\text {cos }}$ | 0.365 | 0.312 | 1.370 | $\mathbf{0 . 2 0 8}$ | $\underline{0.219}$ |
| $\phi_{\exp }$ | 0.120 | 0.610 | 1.000 | 0.147 | 0.140 |
| $\mu_{\text {exp }}$ | 0.086 | 0.309 | 4.631 | 0.103 | 0.204 |
| AVG | $\mathbf{0 . 1 0 3}$ | 0.460 | 2.816 | $\underline{0.125}$ | 0.172 |
| ACTIVE | 2.369 | 2.370 | 1.315 | 2.580 | $\mathbf{2 . 5 9 2}$ |
| SEISMIC | 3.335 | 3.357 | 2.131 | 3.576 | $\mathbf{3 . 5 7 8}$ |

Experimental Details. For Gibbs-Hawkes and EM-Hawkes, we must select parameters of the GP kernel (Equation (3.7), (3.8) and (3.9)). An arbitrary choice of them can lead to poor performance, and to this end, we apply the standard cross validation based on the log-likelihood. We choose the number of basis functions in $[8,16,32,64,128]$ and $a=b$ from $\left[0.2,0.02, \cdots, 2 \times 10^{-8}\right]$. We found that having many
basis functions leads to a high fitting accuracy, but low speed. So, we use 32 basis functions which provides a suitable balance. In terms of kernel parameters $a, b$ of Equation (3.8), we observed that large values return smooth triggering kernels which have a large distance to the ground truth, while small values result in non-smooth predictions which however have small log-likelihoods. As a result, the values $a, b=0.002$ were chosen. 5000 iterations are run to fit each group and first 1000 are ignored (i.e. burned-in).


Figure 3.4: Learned Hawkes triggering kernels using our non-parametric Bayesian approaches. Each red or blue area shows the estimated posterior distributions of $\phi$, while the solid lines indicate the 10, 50 and 90 percentiles. In Figure (a), a synthetic dataset simulated using $\phi_{\exp }(t)$ (in gray) is fit using Gibbs-Hawkes (in red) and EM-Hawkes (in blue); Figure (b) presents learning outcomes on Twitter data in ACTIVE (in red) and SEISMIC (in blue); Figure (c) presents learning outcomes on Twitter data associated with two categories in the ACTIVE set: Music (in red) and Pets \& Animals (in blue).

Results. The top section of Table 3.2 shows the mean relative L2 distance between the learned and the true $\phi$ and $\mu$ on toy data. First, Gibbs-Hawkes and EM-Hawkes are the closest the models to the ground truth cosine model according to the average error values $\left(\mathrm{AVG}_{\mathrm{cos}}\right)$. For the exponential simulation model, both approaches gain the second and the third lowest errors respectively among all methods, and as expected, the parametric Hawkes - which uses an exponential kernel - fits the model best. In contrast, the parametric model retrieves the cosine model worse because of its mismatch with the ground truth model. The learned triggering kernels for $\phi_{\exp }$ and $\phi_{\text {cos }}$ by our approaches are shown in Figure 3.4(a) and Figure A. 1 in the appendix. The ODE-based method performs unsatisfactorily on both simulation settings and it is observed that it performs better on $\left(\mu_{\cos }, \phi_{\text {cos }}\right)$ than on $\left(\mu_{\text {exp }}, \phi_{\exp }\right)$. We explain the second observation as that the regularization of the ODE-based method encourages those triggering kernels that have small L2 norms for their derivatives and the derivative of $\phi_{\exp }(x)$ has a larger norm than that of $\phi_{\cos }(x)$. Notably, tuning the hyper-parameters of the WH method is challenging, and Table 3.2 shows the best result obtained after a rather exhaustive experimentation. We speculate that the overall better performance of our approaches is due to the regularization induced by the prior distributions and less difficult hyper-parameter selection. In addition, we also note that EM-Hawkes always performs better at discovering triggering kernels than Gibbs-Hawkes and this observation also holds on the real-life data. Thus, we
conclude that generating multiple samples per iteration tends to improve modeling of the triggering kernels. In summary, compared with state-of-the-art methods, our approaches achieve better performances for data generated by kernels from several parametric classes; as expected, the parametric models are only effective for data generated from their own class.

Effect of Halpin's Procedure. In Section 3.2.4, we show that using Halpin's procedure reduces the complexity of calculating $p_{i j}$ from quadratic to linear. We now empirically validate this speed-up. To distinguish between quadratic and linear complexity, we compute the ratio between running time and data size, shown in Figure 3.2. The ratio when using Halpin's procedure remains roughly constant as data size increases (the ratio increases linearly without the optimization), which implies that Halpin's procedure renders linear calculation of estimating $p_{i j}$ and sampling branching structures. Later, we show the linear complexity of our method on realworld data.

### 3.4.2 Twitter Diffusion Data

We evaluate the performance of our two proposed approaches on two Twitter datasets that consist of retweet cascades. A retweet cascade contains an original tweet, together with its direct and indirect retweets. Current state of the art diffusion modeling approaches [Zhao et al., 2015; Mishra et al., 2016; Rizoiu et al., 2018] are based on the self-exciting assumption: users get in contact with online content, and then diffuse it to their friends, therefore generating a cascading effect. The two datasets we use have been employed in prior works and they are publicly available:
(i) ACTIVE [Rizoiu et al., 2018] owns 41k retweet cascades, each containing at least 20 (re)tweets with links to Youtube videos. It was collected in 2014 and each Youtube video (and therefore each cascade) is associated with a Youtube category, e.g., Music or News.
(ii) SEISMIC [Zhao et al., 2015] owns 166k randomly sampled retweet cascades, collected in from Oct 7 to Nov 7, 2011. Each cascade contains at least 50 tweets.

Setup. The temporal extent of each cascade is scaled to $[0, \pi]$, and assigned to either training or test data with equal probability. We bundle together groups of 30 cascades of similar size, and we estimate one Hawkes process for each bundle. Unlike for the synthetic dataset, for the retweet cascades dataset there is no true Hawkes process to evaluate against. Instead, we measure using log-likelihood how well the learned model generalizes to the test set. We use the same hyper-parameters values as for the synthetic data. Finally, we follow the prior works on these cascade datasets [Zhao et al., 2015; Rizoiu et al., 2018] by setting the background intensity $\mu$ as 0 , because the cascade datasets contain only the information of triggering relationships.

Fitting Performance. For each dataset, we calculate the log-likelihood per event for each tweet cascade obtained by three baselines and our approaches (Table 3.2). Visibly, our proposed methods consistently outperform baselines, with EM-Hawkes performing slightly better than Gibbs-Hawkes (by $0.5 \%$ for ACTIVE and $0.06 \%$ for

SEISMIC). This seems to indicate that online diffusion is influenced by factors not captured by the parametric kernel, therefore justifying the need to learn the Hawkes kernels non-parametrically. As mentioned in the synthetic data part, the WH-based method has a disadvantage of hard-to-tune hyper-parameters, which leads to the worst performance among all methods.

Scalability. To validate the linear complexity of our method, we record running time per iteration of Gibbs-Hawkes on ACTIVE and SEISMIC in Figure 3.3. The running time rises linearly with the number of points increasing, in line with the theoretical analysis. Linear complexity makes our method scalable and applicable on large datasets.

Interpretation. We show in Figure 3.4(b) and 3.4(c) the learned kernels for information diffusions. We notice that the learned kernels appear to be decaying and long-tailed, in accordance with the prior literature. Figure 3.4(b) shows that the kernel learned on SEISMIC is decaying faster than the kernel learned on ACTIVE. This indicates that non-specific (i.e. random) cascades have a faster decay than video-related cascades, presumably due to the fact that Youtube videos stay longer in the human attention. This connection between the type of content and the speed of the decay seems further confirmed in Figure 3.4(c), where we show the learned kernels for two categories in ACTIVE: Music and Pets $\mathcal{E}$ Animals. Cascades relating to Pets $\mathcal{E}$ Animals have a faster decaying kernel than Music, most likely because Music is an ever-green content.

### 3.5 Summary

In this chapter, we provided the first non-parametric Bayesian inference procedure for the Hawkes process which requires no discretization of the input domain and has an advantage of a linear time complexity. Our method iterates between two steps. First, it samples the branching structure, effectively transforming the Hawkes process into a cluster of Poisson processes. Next, it estimates the Hawkes triggering kernel using a non-parametric Bayesian estimation of the intensity of the cluster Poisson processes. We provide both a full posterior sampler and an EM estimation algorithm based on our ideas. We demonstrated our approach can infer flexible triggering kernels on simulated data. On two large Twitter diffusion datasets, our method outperforms the state-of-the-art in held-out likelihood. Moreover, the learned non-parametric kernel reflects the intuitive longevity of different types of content. The linear complexity of our approach is corroborated on both the synthetic and real problems. The present framework is limited to the univariate unmarked Hawkes process and will be extended to marked multivariate Hawkes process.

The Bayesian Hawkes process model in this chapter is a Gaussian latent model so an advanced Laplace approximation based approach, i.e., the Integrated nested Laplace approximation (INLA) [Rue et al., 2009, 2017], is applicable. INLA is applicable to the case of univariate posterior marginals of latent variables and hyperparameters. More specifically, we could apply INLA in each iteration of Gibbs
sampling by introducing an additional prior distribution for hyper-parameters $\boldsymbol{\theta}$. As per Rue et al. [2009], INLA first approximates indirectly the posterior distribution of $\boldsymbol{\theta}$, say $p(\boldsymbol{\theta} \mid \mathcal{D})$, and then approximates the posterior distribution of $\omega_{i}$ by integrating out $\boldsymbol{\theta}, p\left(\omega_{i} \mid \mathcal{D}\right)=\int p\left(\omega_{i} \mid \boldsymbol{\theta}, \mathcal{D}\right) p(\boldsymbol{\theta} \mid \mathcal{D}) \mathrm{d} \boldsymbol{\theta}$ (for algorithmic details see Rue et al. [2009]). Both steps apply Laplace approximation and the main difference from our Laplace approximation application includes: (1) INLA aims to model univariate marginal distributions while we focus on the joint distribution; (2) INLA's indirect approximation to the posterior $\boldsymbol{\theta}$ can result in non-Gaussian distributions, which enables more flexible modeling. (3) INLA considers a distribution for hyper-parameters, which increases computational burden: e.g. integrating out $\boldsymbol{\theta}$ in computing $p\left(\omega_{i} \mid \mathcal{D}\right)$ often needs numerical integration. We may also apply INLA to the Bayesian Hawkes process without the Gibbs sampling framework, which needs deeper investigation.

In the next chapter, we will introduce a sparse Gaussian process modulated Hawkes process model and propose variational inference for it.

## Variational Inference for Sparse Gaussian Process Modulated Hawkes Process

In this chapter, we propose the sparse Gaussian process modulated Hawkes process which employs a novel variational inference schema, has an advantage of a linear time complexity per iteration and scales to large real-world data. Our method is inspired by the variational Bayesian Poisson process (VBPP) [Lloyd et al., 2015] which provides the Bayesian non-parametric inference only for the whole intensity of the Hawkes process without for its components: the background intensity $\mu$ and the triggering kernel $\phi$. Thus, the VBPP loses the internal reactions between points, and developing the variational Bayesian non-parametric inference for the Hawkes process is non-trivial and more challenging than the VBPP. In this paper, we adapt the VBPP for the Hawkes process and term the new approach the variational Bayesian Hawkes process (VBHP). The structure of this chapter is as follows:
(i) In Section 4.1, we first review a prerequisite of our approach, LBPP.
(ii) In Section 4.2, we introduce a Bayesian non-parametric Hawkes process, which employs a sparse Gaussian process modulated triggering kernel and a Gamma distributed background intensity. We propose a new variational inference schema for such a model which is an EM-like two-step iterative algorithm. Specifically, we employ the branching structure of the Hawkes process so that maximization of the evidence lower bound (ELBO) is tractable. As a result, it introduces extra constraints to the ELBO. To deal with the constraints, we apply the expectation-maximization algorithm to minimize the ELBO.
(iii) In Section 4.3, we contribute a tighter ELBO which improves the fitting performance of our model.
(iv) In Section 4.4, we propose a new acceleration trick based on the finite support assumption of the triggering kernel. The new trick enjoys higher efficiency than prior methods and accelerates the variational inference schema to linear time complexity per iteration.
(v) In Section 4.5, we empirically show that VBHP provides more accurate predictions than state-of-the-art methods on synthetic data and on two large online
diffusion datasets. We validate the linear time complexity and faster convergence of our accelerated variational inference schema compared to the Gibbs sampling method, and the practical utility of our tighter ELBO for model selection, which outperforms the common one in model selection.

### 4.1 Variational Bayesian Poisson Process

Before presenting our approach, we first review Variational Bayesian Poisson process (VBPP) [Lloyd et al., 2015], which is a prerequisite of our method. VBPP applies the VI to the Bayesian Poisson process, which exploits the sparse Gaussian process (GP) to model the Poisson intensity. Specifically, VBPP uses a squared link function to map a sparse GP distributed function $f$ to the Poisson intensity $\lambda(\cdot)=f^{2}(\cdot)$. The sparse GP employs the ARD kernel for given $x, y \in \mathbb{R}^{R}$ :

$$
k(\boldsymbol{x}, \boldsymbol{y}):=\gamma \prod_{r=1}^{R} \exp \left(-\frac{\left(x_{r}-y_{r}\right)^{2}}{2 \alpha_{r}}\right)
$$

where $\gamma$ and $\left\{\alpha_{r}\right\}_{r=1}^{R}$ are GP hyper-parameters. Let $u:=\left(f\left(z_{1}\right), f\left(z_{2}\right), \cdots, f\left(z_{M}\right)\right)$ where $\left\{\boldsymbol{z}_{i}\right\}_{i=1}^{M}$ are inducing points. The prior and the approximate posterior distributions of $u$ are Gaussian distributions,

$$
p(\boldsymbol{u})=\mathcal{N}\left(\boldsymbol{u} \mid \mathbf{0}, \boldsymbol{K}_{z z}\right) \quad q(\boldsymbol{u})=\mathcal{N}(\boldsymbol{u} \mid \boldsymbol{m}, \boldsymbol{S})
$$

where $\boldsymbol{m}$ is the mean vector, and $K_{z z}$ and $S$ are the covariance matrices. Note both $u$ and function evaluations $f$ employ zero mean priors. Notations of VBPP are connected with those of VI (Section 2.2.4) in Table 4.1.

Table 4.1: Notations.

| VBHP | VBPP | VI |
| :---: | :---: | :---: |
| $\mathcal{D}:=\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{N}$ | $\mathcal{D}:=\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{N}$ | $\mathcal{D}$ |
| $B, \mu, f, \boldsymbol{u}$ | $f, \boldsymbol{u}$ | $f$ |
| $k_{0}, c_{0},\left\{\alpha_{i}\right\}_{i=1}^{R}, \gamma$ | $\left\{\alpha_{i}\right\}_{i=1}^{R}, \gamma$ | $\theta$ |
| $k, c, \boldsymbol{m}, \boldsymbol{S},\left\{\alpha_{i}\right\}_{i=1}^{R}, \gamma,\left\{\left\{q_{i j}\right\}_{j=0}^{i-1}\right\}_{i=1}^{N}$ | $\boldsymbol{m}, \boldsymbol{S},\left\{\alpha_{i}\right\}_{i=1}^{R}, \gamma$ | $\theta^{\prime}$ |

Importantly, the variational joint distribution of $f$ and $u$ uses the exact conditional distribution $p(f \mid \boldsymbol{u})$, i.e.,

$$
\begin{equation*}
q(f, \boldsymbol{u}) \equiv p(f \mid \boldsymbol{u}) q(\boldsymbol{u}) \tag{4.1}
\end{equation*}
$$

which in turn leads to the posterior GP:

$$
\begin{align*}
q(f) & =\mathcal{N}(f \mid v, \Sigma),  \tag{4.2}\\
v(\boldsymbol{x}) & \equiv \boldsymbol{K}_{x z} \boldsymbol{K}_{z z}^{-1} \boldsymbol{m}, \\
\Sigma\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right) & \equiv \boldsymbol{K}_{x x^{\prime}}+\boldsymbol{K}_{x z} \boldsymbol{K}_{z z}^{-1}\left(\boldsymbol{S} \boldsymbol{K}_{z z}^{-1}-I\right) \boldsymbol{K}_{z x^{\prime}} .
\end{align*}
$$

Then, the ELBO is obtained by using Equation (2.7):

$$
\operatorname{ELBO}(q(f, \boldsymbol{u}), p(\mathcal{D} \mid f, \boldsymbol{u}), p(f, \boldsymbol{u}))=\mathbb{E}_{q(f)}[\log p(\mathcal{D} \mid f)]-\operatorname{KL}(q(\boldsymbol{u}) \| p(\boldsymbol{u}))
$$

Note that the second term is the KL divergence between two multivariate Gaussian distributions, so is available in closed form. The first term turns out to be the expectation w.r.t. $q(f)$ of the $\log$-likelihood $\log p(\mathcal{D} \mid f)=\sum_{i=1}^{N} \log f^{2}\left(x_{i}\right)-\int_{\mathcal{T}} f^{2}$. The expectation of the integral part is relatively straight-forward to compute and the expectation of the other (data-dependent) part is available in almost closed-form with a hyper-geometric function.

### 4.2 Variational Bayesian Hawkes Process

### 4.2.1 Notations

To extend VBPP to Hawkes process, we introduce two more variables: the background intensity $\mu$ and the branching structure $B$, defined in Section 2.1.2. We assume that the prior distribution of $\mu$ is a Gamma distribution $p(\mu)=\operatorname{Gamma}\left(\mu \mid k_{0}, c_{0}\right)^{1}$ and the posterior distribution is approximated by another Gamma distribution $q(\mu)=$ $\operatorname{Gamma}(\mu \mid k, c)$. Given a dataset $\mathcal{D}=\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{N}$, the branching structure is expressed as $B=\left\{\boldsymbol{b}_{i}\right\}_{i=1}^{N}$ with $\boldsymbol{b}_{i}=\left\{b_{i j}\right\}_{j=0}^{i-1}$ defined around Equation (2.1). As $b_{i j} \in\{0,1\}$ is a binary variable and there is only one $b_{i j}$ equal to 1 among $\boldsymbol{b}_{i}=\left\{b_{i j}\right\}_{j=0}^{i-1}$ (one point has a unique parent), $\boldsymbol{b}_{i}$ hence has a categorical distribution. We let the variational posterior probability of $b_{i j}=1$ be denoted as $q_{i j}=q\left(b_{i j}=1\right)$ and there must be $\sum_{j=0}^{i-1} q_{i j}=1$. As a result, the variational posterior probability of $B$ is expressed as:

$$
\begin{equation*}
q(B)=\prod_{i=1}^{N} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}} \tag{4.3}
\end{equation*}
$$

The same squared link function is adopted for the triggering kernel $\phi(\cdot)=f^{2}(\cdot)$, so are the priors for $f$ and $\boldsymbol{u}$, namely $\mathcal{N}\left(f \mid \mathbf{0}, \mathbf{K}_{x x^{\prime}}\right)$ and $\mathcal{N}\left(\boldsymbol{u} \mid \mathbf{0}, \mathbf{K}_{z z^{\prime}}\right)$. More link functions such as $\exp (\cdot)$ are discussed by Lloyd et al. [2015]. Moreover, we use the same variational joint posterior on $f$ and $u$ as Equation (4.1). Consequently, we complete the variational joint distribution on all latent variables as below:

$$
\begin{equation*}
q(B, \mu, f, \boldsymbol{u}) \equiv q(B) q(\mu) p(f \mid \boldsymbol{u}) q(\boldsymbol{u}), \tag{4.4}
\end{equation*}
$$

and notations of VBHP are summarized in Table 4.1. The variational posterior probability of $B$ (Equation (4.3)) has the same form as that of the true posterior, so it is a good approximation. The variational joint distribution $q(B, \mu, f, u)$ assumes independence between the modeled variables, which is not the case for the true joint posterior distribution, e.g. the branching structure $B$ is closely dependent on the

[^0]background intensity $\mu$ and the function $f$. Thus such an approximation is not perfect. The approximation error is insignificant in practice and our experiments demonstrate the effectiveness of the proposed approximation.

Based on Equation (2.7) and (4.4), we obtain the ELBO for VBHP as below (see details in Section B. 1 of the appendix). To differentiate with the tighter ELBO presented later in Section 4.3, we name the below one as the common ELBO (CELBO).

$$
\begin{align*}
& \operatorname{CELBO}(q(B, \mu, f, \boldsymbol{u}), p(\mathcal{D} \mid B, \mu, f, \boldsymbol{u}), p(B, \mu, f, \boldsymbol{u})) \\
& =\underbrace{\mathbb{E}_{q(B, \mu, f)}[\log p(\mathcal{D}, B \mid f, \mu)]}_{\text {Data Dependent Expectation (DDE) }}+H_{B}-\operatorname{KL}(q(\mu) \| p(\mu))-\operatorname{KL}(q(\boldsymbol{u}) \| p(\boldsymbol{u})) . \tag{4.5}
\end{align*}
$$

where $H_{B}=-\sum_{i=1}^{N} \sum_{j=0}^{i-1} q_{i j} \log q_{i j}$ is the entropy of the variational posterior $B$. The KL terms are between gamma and Gaussian distributions, for which closed forms are provided in Section B. 2 of the appendix.

### 4.2.2 Data Dependent Expectation

Now, we are left with the problem of computing the data dependent expectation (DDE) in Equation (4.5). The DDE is w.r.t. the variational posterior probability $q(B, \mu, f)$. From Equation (4.4), $q(B, \mu, f)=\int q(B, \mu, f, \mathbf{u}) \mathrm{d} \mathbf{u}=q(B) q(\mu) q(f)$ and $q(f)$ is identical to Equation (4.2). As a result, we can compute the DDE w.r.t. $q(B)$ first, and then w.r.t. $q(\mu)$ and $q(f)$.

Expectation w.r.t. $q(B)$. From Equation (2.3), we easily obtain $\log p(\mathcal{D}, B \mid f, \mu)$ by replacing $\phi$ with $f^{2}$, whereupon it is clear that only $b_{i j}$ in $\log p(\mathcal{D}, B \mid f, \mu)$ is dependent on $B$. Therefore, $\mathbb{E}_{q(B)}[\log p(\mathcal{D}, B \mid f, \mu)]$ is computed as:

$$
\mathbb{E}_{q(B)}[\log p(\mathcal{D}, B \mid f, \mu)]=\sum_{i=1}^{N}\left(\sum_{j=1}^{i-1} q_{i j} \log f_{i j}^{2}+q_{i 0} \log \mu\right)-\sum_{i=1}^{N} \int_{\mathcal{T}_{i}} f^{2}-\mu|\mathcal{T}|
$$

where $f_{i j}:=f\left(x_{i}-\boldsymbol{x}_{j}\right)$.
Expectation w.r.t. $q(f)$ and $q(\mu)$. We compute the expectation w.r.t. $q(f)$ and $q(\mu)$ by exploiting the expectation and the log expectation of the Gamma distribution: $\mathbb{E}_{q(\mu)}(\mu)=k c$ and $\mathbb{E}_{q(\mu)}(\log (\mu))=\psi(k)+\log c$, and also the property $\mathbb{E}\left(x^{2}\right)=$ $\mathbb{E}(x)^{2}+\operatorname{Var}(x):$

$$
\begin{aligned}
\mathrm{DDE}= & \sum_{i=1}^{N}\left[\sum_{j=1}^{i-1} q_{i j} \mathbb{E}_{q(f)}\left(\log f_{i j}^{2}\right)+q_{i 0}(\psi(k)+\log c)\right. \\
& \left.-\int_{\mathcal{T}_{i}} \mathbb{E}_{q(f)}^{2}(f)-\int_{\mathcal{T}_{i}} \operatorname{Var}_{q(f)}(f)\right]-k c|\mathcal{T}|
\end{aligned}
$$

where $\psi$ is the Digamma function. We provide closed form expressions for $\int_{\mathcal{T}_{i}} \mathbb{E}_{q(f)}^{2}(f)$
and $\int_{\mathcal{T}_{i}} \operatorname{Var}_{q(f)}(f)$ in Section B. 3 of the appendix. As in VBPP,

$$
\mathbb{E}_{q(f)}\left(\log f_{i j}^{2}\right)=-\widetilde{G}\left(-v_{i j}^{2} /\left(2 \Sigma_{i j}\right)\right)+\log \left(\Sigma_{i j} / 2\right)-C
$$

is available in the closed form with a hyper-geometric function, where $v_{i j}=v\left(x_{i}-\boldsymbol{x}_{j}\right)$, $\Sigma_{i j}=\Sigma\left(x_{i}-x_{j}, x_{i}-x_{j}\right)$ ( $v$ and $\Sigma$ are defined in Equation (4.2)), $C \approx 0.57721566$ is the Euler-Mascheroni constant and $\widetilde{G}$ is defined as

$$
\widetilde{G}(z)={ }_{1} F_{1}^{(1,0,0)}(0,1 / 2, z),
$$

that is, the partial derivative of the confluent hyper-geometric function ${ }_{1} F_{1}$ w.r.t. the first argument. We compute $\widetilde{G}$ using the method of Ancarani and Gasaneo [2008] and implement $\tilde{G}$ and $\tilde{G}^{\prime}$ by linear interpolation of a lookup table.

### 4.2.3 Predictive Distribution of $\phi$

The predictive distribution of $f(x)$ depends on the posterior $\mathbf{u}$. We assume that the optimal variational distribution of $\mathbf{u}$ approximates the true posterior distribution, namely $q\left(\mathbf{u} \mid \mathcal{D}, \boldsymbol{\theta}^{\prime *}\right)=\mathcal{N}\left(\mathbf{u} \mid \boldsymbol{m}^{*}, \mathbf{S}^{*}\right) \approx p(\mathbf{u} \mid \mathcal{D}, \boldsymbol{\theta})$. Therefore, there is

$$
\begin{aligned}
q\left(f \mid \mathcal{D}, \boldsymbol{\theta}^{\prime *}\right) & =\int p(f \mid \boldsymbol{u}) q\left(\mathbf{u} \mid \mathcal{D}, \boldsymbol{\theta}^{\prime *}\right) \mathrm{d} \boldsymbol{\theta} \\
& \approx \int p(f \mid \boldsymbol{u}) p(\mathbf{u} \mid \mathcal{D}, \boldsymbol{\theta}) \mathrm{d} \boldsymbol{\theta}=p(f \mid \mathcal{D}, \boldsymbol{\theta}),
\end{aligned}
$$

and we thus use $q\left(f(\tilde{x}) \mid \mathcal{D}, \boldsymbol{\theta}^{\prime *}\right)$ as the the approximate predictive distribution which is calculated as

$$
\begin{aligned}
q\left(f(\tilde{\boldsymbol{x}}) \mid \mathcal{D}, \boldsymbol{\theta}^{\prime *}\right) & =\int p(f(\tilde{\boldsymbol{x}}) \mid \boldsymbol{u}) q\left(\boldsymbol{u} \mid \mathcal{D}, \boldsymbol{\theta}^{\prime *}\right) \mathrm{d} \boldsymbol{u} \\
& =\mathcal{N}\left(\boldsymbol{K}_{\tilde{x} z} \boldsymbol{K}_{z z}^{-1} \boldsymbol{m}^{*}, \boldsymbol{K}_{\tilde{x} \tilde{x}}-\boldsymbol{K}_{\tilde{x} z} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{K}_{z \tilde{x}}+\boldsymbol{K}_{\tilde{x} z} \boldsymbol{K}_{z z^{\prime}}^{-1} \mathbf{S}^{*} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{K}_{z \tilde{x}}\right) \\
& \equiv \mathcal{N}\left(\tilde{v}, \tilde{\sigma}^{2}\right) .
\end{aligned}
$$

Given the relation $\phi=f^{2}$, it is straightforward to derive the corresponding approximate posterior of $\phi(\tilde{x})$

$$
\phi(\tilde{x}) \sim \operatorname{Gamma}(\tilde{k}, \tilde{c})
$$

where the shape $\tilde{k}=\left(\tilde{v}^{2}+\tilde{\sigma}^{2}\right)^{2} /\left[2 \tilde{\sigma}^{2}\left(2 \tilde{v}^{2}+\tilde{\sigma}^{2}\right)\right]$ and the scale $\tilde{c}=2 \tilde{\sigma}^{2}\left(2 \tilde{v}^{2}+\tilde{\sigma}^{2}\right) /\left(\tilde{v}^{2}+\right.$ $\tilde{\sigma}^{2}$ ).

### 4.3 New Variational Inference Schema

We now propose a new variational inference (VI) schema which uses a tighter ELBO than the common one, i.e.

Theorem 1. For VBHP, there is a tighter ELBO

$$
\underbrace{\mathbb{E}_{q(B, \mu, f)}[\log p(\mathcal{D}, B \mid f, \mu)]+H_{B}}_{\equiv \operatorname{TELBO}} \leq \log p(\mathcal{D}) .
$$

Remark. TELBO is tighter because it is equivalent to the CELBO (Equation (4.5)) except without subtracting non-negative KL divergences over $\mu$ and $\mathbf{u}$. Thus, it is easy to see that the gap between the TELBO and $\log p(\mathcal{D})$ is
$\log p(\mathcal{D})-\operatorname{TELBO}=-K L(q(\mu) \| p(\mu))-K L(q(\boldsymbol{u}) \| p(\boldsymbol{u}))+K L(q(B, \mu, f, \boldsymbol{u}) \| p(B, \mu, f, \boldsymbol{u} \mid \mathcal{D}))$.
The aggregate effect of the three KL terms is rather challenging to understand and needs further investigation. Other graphical models such as the variational Gaussian mixture model [Attias, 1999] have a similar TELBO. Later on, we propose a new VI schema based on the TELBO, where the TELBO will be applied to selecting the hyper-parameters as it provides a tighter approximation to the log marginal likelihood compared to CELBO.

Proof. With the variational posterior probability of the branching structure $q(B)$ defined in Equation (4.3) and through the Jensen's inequality, we have:

$$
\begin{equation*}
\log p(\mathcal{D}) \geq \sum_{B} q(B) \log p(\mathcal{D}, B)+H_{B}, \tag{4.6}
\end{equation*}
$$

where $H_{B}$ is the entropy of $B$ defined in Equation (4.5). The term $\sum_{B} q(B) \log p(\mathcal{D}, B)$ can be understood as follows. Consider that infinite branching structures are drawn from $q(B)$ independently, say $\left\{B_{i}\right\}_{i=1}^{\infty}$. Given a branching structure $B_{i}$, the Hawkes process can be decomposed into a cluster of Poisson processes, denoted as ( $\mathcal{D}, B_{i}$ ), and the corresponding $\log$-likelihood is $\log p\left(\mathcal{D}, B_{i}\right)$. Then, $\sum_{B} q(B) \log p(\mathcal{D}, B)$ is the mean of all $\log$ likelihoods $\left\{\log p\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{\infty}$,

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^{n} \log p\left(\mathcal{D}, B_{i}\right)=\lim _{n \rightarrow \infty} \sum_{B} \frac{n_{B}}{n} \log p(\mathcal{D}, B)=\sum_{B} q(B) \log p(\mathcal{D}, B), \tag{4.7}
\end{equation*}
$$

where $n_{B}$ is the number of occurrences of branching structure $B$. Since all branching structures $\left\{B_{i}\right\}_{i=1}^{\infty}$ are i.i.d., the clusters of Poisson processes generated over $\left\{B_{i}\right\}_{i=1}^{\infty}$ should also be independent, i.e., $\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{\infty}$ are i.i.d.. It follows that

$$
\begin{equation*}
\sum_{i=1}^{\infty} \log p\left(\mathcal{D}, B_{i}\right)=\log p\left(\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{\infty}\right) . \tag{4.8}
\end{equation*}
$$

We compute the CELBO of $\log p\left(\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{\infty}\right)$ by making $z=(\mu, f, \mathbf{u})$ and $\boldsymbol{x}=$ $\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{n}$ in Equation (2.7):

$$
\begin{gather*}
\log p\left(\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{n}\right) \geq E_{q(f, \mu)}\left[\log p\left(\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{n} \mid f, \mu\right)\right]  \tag{4.9}\\
-\operatorname{KL}(q(\mu) \| p(\mu))-\operatorname{KL}(q(\boldsymbol{u}) \| p(\boldsymbol{u})) .
\end{gather*}
$$

Further, we plug Equation (4.8) and (4.9) into Equation (4.7):

$$
\begin{aligned}
& \text { Equation (4.7) }= \lim _{n \rightarrow \infty} \frac{1}{n} \log p\left(\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{n}\right) \\
& \stackrel{(a)}{\geq} \\
& \lim _{n \rightarrow \infty} \frac{1}{n} E_{q(f, \mu)}\left[\log p\left(\left\{\left(\mathcal{D}, B_{i}\right)\right\}_{i=1}^{n} \mid f, \mu\right)\right] \\
& \stackrel{(b)}{=} \lim _{n \rightarrow \infty} \sum_{n_{B}} \frac{n_{B}}{n} E_{q(f, \mu)}[\log p(\mathcal{D}, B \mid f, \mu)] \\
&=E_{q(f, \mu, B)}[\log p(\mathcal{D}, B \mid f, \mu)]
\end{aligned}
$$

where (a) is because the finite values of KL terms are divided by infinitely large $n$, and (b) is due to i.i.d. ( $\mathcal{D}, B_{i}$ ) and the variational posterior $B$ being independent of $f$ and $\mu$. Finally, we plug the above inequality into Equation (4.6) and obtain the TELBO.

### 4.3.1 New Optimization Schema for VBHP

To optimize the model parameters under constraints $\sum_{j=0}^{i-1} q_{i j}=1$, we employ the expectation-maximization algorithm. Specifically, in the E step, all $q_{i j}$ are optimized to maximize the CELBO, and in the $\mathbf{M}$ step, $m, S, k$ and $c$ are updated to increase the CELBO. We don't use the TELBO to optimize the variational distributions because it doesn't guarantee minimizing the KL divergence between variational and true posterior distributions. Instead, the TELBO is employed to select GP hyper-parameters:

$$
\left\{\alpha_{i}^{*}\right\}_{i=1}^{R}, \gamma^{*}=\underset{\left\{\alpha_{i}\right\}_{i=1}^{R}, \gamma}{\operatorname{argmax}} \text { TELBO. }
$$

The TELBO bounds the marginal likelihood more tightly than CELBO, and is therefore expected to lead to a better predictive performance - an intuition which we empirically validate in Section 4.5.

The updating equations for $q_{i j}$ are derived through maximization of Equation (4.5) under the constraints $\sum_{j=0}^{i-1} q_{i j}=1$ for all $i$. This maximization problem is dealt with the Lagrange multiplier method, and yields the below updating equations:

$$
q_{i j}=\left\{\begin{array}{l}
\exp \left(\mathbb{E}_{q(f)}\left(\log f_{i j}^{2}\right)\right) / A_{i}, \quad j>0 ; \\
\theta \exp (\psi(k)) / A_{i}, \quad j=0,
\end{array}\right.
$$

where $A_{i}=\theta \exp (\psi(k))+\sum_{j=1}^{i-1} \exp \left(\mathbb{E}_{q(f)}\left(\log f_{i j}^{2}\right)\right)$ is the normalizer.
Furthermore, and similarly to VBPP, we fix the inducing points on a regular grid over $\mathcal{T}$. Despite the observation that more inducing points lead to better fitting accuracy [Lloyd et al., 2015; Snelson and Ghahramani, 2006], in the case of our more complex VBHP, more inducing points may cause slow convergence (Figure B.2(a) in the appendix) for some hyper-parameters, and therefore lead to poor performance in limited iterations. Generally, more inducing points improve accuracy at the expense of longer fitting time.

### 4.4 Acceleration Trick

### 4.4.1 Time Complexity Without Acceleration

In the E step of model optimization, updating $q_{i j}$ requires computing the mean and the variance of all $f_{i j}$, which both take $O\left(M^{3}+M^{2} N^{2}\right)$ with $N$ points in the HP and $M$ inducing points. Here, we omit the dimension of data $R$ since normally $M>R$ for a regular grid of inducing points. Similarly, in the $M$ step, computing the hypergeometric term requires the means and variances of all the $f_{i j}$. Finally, computation of the integral terms takes $O\left(M^{3} N\right)$. Thus, the total time complexity per iteration is $O\left(M^{3} N+M^{2} N^{2}\right)$.

### 4.4.2 Acceleration to Linear Time Complexity

To accelerate our VBHP, similarly to Zhang et al. [2019] we exploit the finite support assumption of the triggering kernel, assuming the kernel has negligible values for sufficiently large inputs. As a result, sufficiently distant pairs of points do not enter into the computations. This trick reduces possible parents of a point from all prior points to a set of neighbors. The number of relevant neighbors is bounded by a constant $C$ and as a result the total time complexity is reduced to $O\left(C M^{3} N\right)$.

Specifically, we introduce a compact region $\mathcal{S}=\times_{r=1}^{R}\left[\mathcal{S}_{r}^{\min }, \mathcal{S}_{r}^{\max }\right] \subseteq \mathcal{T}$ so that $\phi\left(x_{i}-x_{j}\right)=0$ and $q_{i j}=0$ if $x_{i}-x_{j} \notin \mathcal{S}$. As a result, all terms related to $x_{i}-x_{j} \notin \mathcal{S}$ vanish. To choose a suitable $\mathcal{S}$, we again use the TELBO, taking the smallest $S$ for which the TELBO doesn't drop significantly; we optimize $S_{r}^{\min }$ and $S_{r}^{\max }$ by grid search with other dimensions fixed (so that this step is run $R$ times in total) and we optimize $S_{r}^{\min }$ after optimizing $S_{r}^{\text {max }}$.

Rather than selecting pairs of points in each iteration in the manner of Halpin's trick [Halpin, 2012; Zhang et al., 2019], our method pre-computes those pairs, leading to gains in computational efficiency. The similar aspect is that both tricks have hyperparameters to select to threshold the triggering kernel value. We employ the TELBO for hyper-parameter selection while frequentist methods use the cross validation.

### 4.5 Experiments

### 4.5.1 Evaluation

We employ two metrics: the first is the $\mathbf{L}_{\mathbf{2}}$ distance (for cases with a known ground truth), which measures the difference between predictive and truth Hawkes kernels, formulated as $\mathrm{L}_{2}\left(\phi_{\text {pred }}, \phi_{\text {true }}\right)=\left(\int_{\mathcal{T}}\left(\phi_{\text {pred }}(x)-\phi_{\text {true }}(x)\right)^{2} \mathrm{~d} x\right)^{0.5}$ and $\mathrm{L}_{2}\left(\mu_{\text {pred }}, \mu_{\text {true }}\right)=$ $\left|\mu_{\text {pred }}-\mu_{\text {true }}\right|$; the second is the hold-out log likelihood (HLL), which describes how well the predictive model fits the test data, formulated as $\log p\left(\mathcal{D}_{\text {Test }}=\left\{x_{i}\right\}_{i=1}^{N} \mid \mu, f\right)=$ $\sum_{i=1}^{N} \log \lambda\left(x_{i}\right)-\int_{\mathcal{T}} \lambda$. To calculate the HLL for each process, we generate a number of test sequences by every time randomly assigning each point of the original process to either a training or testing sequence with equal probability; HLLs of test sequences are normalized (by dividing test sequence length) and averaged.


Figure 4.1: The relationship between the $\log$ marginal likelihood and the $L_{2}$ distance. In (a), the true $\phi_{\sin }$ (dash green) is plotted with the median (solid) and the [0.1, 0.9] interval (filled) of the approximate posterior triggering kernel obtained by VBHP and Gibbs Hawkes (10 inducing points). It uses the maximum point of the TELBO (red star in (b)). In (c), the maximum point of the TELBO is marked. The maximum point overlaps with that of the CELBO. $[0,1.4]$ is used as the support of the predictive triggering kernel and 10 inducing points are used.

### 4.5.2 Prediction

We use the pointwise mode of the approximate posterior triggering kernel as the prediction because it is computationally intractable to find the posterior mode at multiple point locations [Zhang et al., 2019]. Besides, we exploit the mode of the approximate posterior background intensity as the predictive background intensity.

### 4.5.3 Baselines

We use the following models as baselines.
(i) A parametric Hawkes process equipped with the sum of exponential (SumExp) triggering kernel $\phi(x)=\sum_{i=1}^{K} a_{1}^{i} a_{2}^{i} \exp \left(-a_{2}^{i} x\right)$ and the constant background intensity.
(ii) The ordinary differential equation (ODE) based non-parametric non-Bayesian Hawkes process [Zhou et al., 2013]. The code is publicly available [Bacry et al., 2017].
(iii) Wiener-Hopf (WH) equation based non-parametric non-Bayesian Hawkes process [Zhou et al., 2013]. The code is publicly available [Bacry et al., 2017].
(iv) The Gibbs sampling based Bayesian non-parametric Hawkes process (Gibbs Hawkes) [Zhang et al., 2019].

For fairness, the ARD kernel is used by Gibbs Hawkes and corresponding eigenfunctions are approximated by Nyström method [Williams and Seeger, 2001], where regular grid points are used as VBHP. Different from batch training in [Zhang et al., 2019], all experiments are conducted on single sequences.


Figure 4.2: Figure (a), (b): The relationship between the TELBO and the HLL. Figure (c), (d): Average fitting time (seconds) per iteration. In Figure (a), the maximum point is marked by the red star. In Figure (b), the maximum points of the TELBO and CELBO are marked by red and blue stars. Figure (c) is plotted on 50 processes. Figure (d) shows the fitting time of Gibbs Hawkes (star) and VBHP (circle) on 120 processes. 10 inducing points are used unless specified.

Table 4.2: Results on synthetic and real-world data (mean $\pm$ one standard variance). VBHP (C) and (T) use the CELBO and the TELBO to update the hyper-parameters respectively. Bold numbers denote the best performance.

| Measure | Data | SumExp | ODE | WH | Gibbs Hawkes | VBHP (C) | VBHP ( T ) |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\mathrm{L}_{2}$ | Sin | $\phi: 0.693_{ \pm 0.028}$ | $0.665_{ \pm 0.121}$ | $2.463_{ \pm 0.145}$ | $0.408_{ \pm 0.198}$ | $\mathbf{0 . 1 5 2}_{ \pm 0.091}$ | $0.183_{ \pm 0.076}$ |
|  |  | $\mu: 2.968 \pm 1.640$ | $4.514_{ \pm 3.808}$ | $6.794_{ \pm 5.054}$ | $4.108_{ \pm 3.949}$ | $0.640_{ \pm 0.528}$ | $\mathbf{0 . 5 7 9} \pm 0.523$ |
|  | Cos | $\phi: 0.473_{ \pm 0.102}$ | $0.697 \pm 0.065$ | $1.743_{ \pm 0.083}$ | $0.667 \pm 0.686$ | $0.325_{ \pm 0.073}$ | $\mathbf{0 . 2 9 2} \pm 0.096$ |
|  |  | $\mu: 2.751_{ \pm 1.902}$ | $7.030_{ \pm 5.662}$ | $6.099_{ \pm 4.613}$ | $4.685{ }_{ \pm 4.421}$ | $0.555_{ \pm 0.294}$ | $\mathbf{0 . 5 1 5} \pm 0.293$ |
|  | Exp | $\phi: \mathbf{0 . 1 3 3}_{ \pm 0.138}$ | $1.835_{ \pm 0.539}$ | $2.254_{ \pm 2.042}$ | $0.676_{ \pm 0.233}$ | $0.257_{ \pm 0.086}$ | $0.235_{ \pm 0.102}$ |
|  |  | $\mu: 3.290_{ \pm 1.991}$ | $8.969_{ \pm 8.604}$ | $16.66 \pm 20.95$ | $7.648_{ \pm 9.647}$ | $\mathbf{0 . 4 7 1}{ }_{ \pm 0.432}$ | $0.486_{ \pm 0.418}$ |
| HLL | Sin | $3.490_{ \pm 0.400}$ | $3.489_{ \pm 0.413}$ | $3.233_{ \pm 0.273}$ | $3.492_{ \pm 0.406}$ | $3.488_{ \pm 0.400}$ | $3.497 \pm 0.406$ |
|  | Cos | $3.874_{ \pm 0.544}$ | $3.872_{ \pm 0.552}$ | $3.613_{ \pm 0.373}$ | $3.871_{ \pm 0.562}$ | $3.876_{ \pm 0.541}$ | $\mathbf{3 . 8 7 8} \pm 0.548$ |
|  | Exp | $2.825_{ \pm 0.481}$ | $2.822_{ \pm 0.496}$ | $2.782_{ \pm 0.490}$ | $2.826_{ \pm 0.492}$ | $2.826 \pm 0.491$ | $2.829 \pm 0.487$ |
|  | ACTIVE | $1.692_{ \pm 1.371}$ | $0.880_{ \pm 2.716}$ | $0.710_{ \pm 0.943}$ | $1.323_{ \pm 2.160}$ | $1.824_{ \pm 1.159}$ | $1.867 \pm 1.181$ |
|  | SEISMIC | $2.943_{ \pm 0.959}$ | $2.582_{ \pm 1.665}$ | $1.489_{ \pm 1.796}$ | $3.110_{ \pm 1.251}$ | $3.143_{ \pm 0.895}$ | $\mathbf{3 . 1 6 4} \pm 0.843$ |

### 4.5.4 Synthetic Experiments

Synthetic Data. Our synthetic data are generated from three Hawkes processes over $\mathcal{T}=[0, \pi]$, whose triggering kernels are sin, cos and exp functions respectively, shown as below, and whose background intensities are the same $\mu=10$ :

$$
\begin{aligned}
& \phi_{\sin }(x)=0.9[\sin (3 x)+1], x \in[0, \pi / 2] ; \text { otherwise, } 0 ; \\
& \phi_{\cos }(x)=\cos (2 x)+1, x \in[0, \pi / 2] ; \text { otherwise, } 0 ; \\
& \phi_{\exp }(x)=5 \exp (-5 x), x \in[0, \infty) .
\end{aligned}
$$

As a result, for any generated sequence, say $\left\{x_{i}\right\}_{i=1}^{N}, \mathcal{T}_{i}=\left[0, \pi-x_{i}\right]$ is used in the CELBO and the TELBO. We can check that all three triggering kernels have negligible values when the input is large, so the finite support assumption is satisfied.

Model Selection. As the marginal likelihood $p(\mathcal{D} \mid \boldsymbol{\theta})$ is a key advantage of our method over non-Bayesian approaches [Zhou et al., 2013; Bacry and Muzy, 2016], we investigate its efficacy for model selection. Figure 4.1(b) shows the contour plot of the approximate log marginal likelihood (the TELBO) of a sequence. It is observed that the contour plot of the TELBO has agreement to the contour plots of $\mathrm{L}_{2}(\phi)$ (Figure 4.1(c)) GP hyper-parameters with relatively high marginal likelihoods have relatively low $\mathrm{L}_{2}$ errors. Figure 4.1(a) plots the posterior triggering kernel corresponding to the maximal approximate marginal likelihood. Similar agreement is also observed between the TELBO and the HLL (Figure 4.2(a), 4.2(b)). This demonstrates the practical utility of both the marginal likelihood itself and our approximation of it.

Evaluation. To evaluate VBHP on synthetic data, 20 sequences are drawn from each model and 100 pairs of train and test sequences drawn from each sample to compute the HLL. We select GP hyper-parameters of Gibbs Hawkes and of VBHP by maximizing approximate marginal likelihoods. Table 4.2 shows evaluations for baselines and VBHP (using 10 inducing points for trade-off between accuracy and time, so does Gibbs Hawkes) in both $L_{2}$ and HLL. VBHP achieves the best performance but is two orders of magnitudes slower than Gibbs Hawkes per iteration (shown as Figure 4.2(c) and 4.2(d)). The speed of VBHP is limited by its complicated implementation, such as the linear interpolation of the lookup table when computing $\tilde{G}$ as Section 4.2.2, and we could accelerate it with more advanced implementation techniques. Although the Gibbs Hawkes is based on the Markov Chain Monte Carlo algorithm and is expected to return more accurate results, it employs Laplace approximation per iteration and leads to approximation error, which is one of the causes of the inferior performance. The TELBO performs closely to the CELBO in the $\mathrm{L}_{2}$ error and this is also reflected in Figure 4.1(c) where the maximum points of the TELBO and the CELBO overlap. In contrast, the TELBO consistently improves the performance of VBHP in the HLL, which is also reflected in Figure 4.2(b) where hyper-parameters selected by the TELBO tend to have a higher HLL. Interestingly, when the parametric model SumExp uses the same triggering kernel (a single exponential function) as the ground truth $\phi_{\text {exp }}$, SumExp fits $\phi_{\exp }$ best in $L_{2}$ distance while due to learning on single sequences, the background intensity has relatively high errors. Although our method
is not aware of the parametric family of the ground truth, it performs well. Compared with non-parametric frequentist methods which have strong fitting capacity but suffer from noisy data and have difficulties with hyper-parameter selection, our Bayesian solution overcomes these disadvantages and achieves better performance.

### 4.5.5 Real World Experiments

Real World Data. We conclude our experiments with two large scale tweet datasets. ACTIVE [Rizoiu et al., 2018] is a tweet dataset which was collected in 2014 and contains $\sim 41 \mathrm{k}(\mathrm{re})$ tweet temporal point processes with links to Youtube videos. Each sequence contains at least 20 (re)tweets. SEISMIC [Zhao et al., 2015] is a large scale tweet dataset which was collected from October 7 to November 7, 2011, and contains $\sim 166 \mathrm{k}$ (re)tweet temporal point processes. Each sequence contains at least 50 (re)tweets.

Evaluation. Similarly to synthetic experiments, we evaluate the fitting performance by averaging HLL of 20 test sequences randomly drawn from each original datum. We scale all original data to $\mathcal{T}=[0, \pi]$ (leading to $\mathcal{T}_{i}=\left[0, \pi-x_{i}\right]$ used in the CELBO and the TELBO for a sequence $\left\{x_{i}\right\}_{i=1}^{N}$ ) and employ 10 inducing points to balance time and accuracy. The model selection is performed by maximizing the approximate marginal likelihood. The obtained results are shown in Table 4.2. Again, we observe similar predictive performance of VBHP: the TELBO performs better the CELBO; VBHP achieves best scores. This demonstrates our Bayesian model and novel VI schema are useful for flexible real life data.

Fitting Time. We further evaluate the fitting speed ${ }^{2}$ of VBHP and Gibbs Hawkes on synthetic and real-world point processes, which is summarized in Figure 4.2(c) and 4.2(d). The fitting time is averaged over iterations and we observe that the increasing trends with the number of inducing points and with the data size are similar between Gibbs Hawkes and VBHP. Although VBHP is significantly slower than Gibbs Hawkes per iteration, VBHP converges faster, in 10~20 iterations (Figure B. 2 of the appendix), giving an average convergence time of 549 seconds for a sequence of 1000 events, compared to 699 seconds for Gibbs Hawkes. The slope of VBHP in Figure 4.2(d) is 1.04 (log-scale) and the correlation coefficient is 0.96 , so we conclude that the fitting time is linear to the data size.

### 4.6 Conclusions

In this chapter, we presented a new Bayesian non-parametric Hawkes process whose triggering kernel is modulated by a sparse Gaussian process and background intensity is Gamma distributed. We provided a novel VI schema for such a model: we employed the branching structure so that the common ELBO is maximized by the expectationmaximization algorithm; we contributed a tighter ELBO which performs better in model selection than the common one. To address the difficulty with scaling with the

[^1]data size, we utilize the finite support assumption of the triggering kernel to reduce the number of possible parents for each point. Different from prior acceleration methods, ours enjoys higher efficiency. On synthetic data and two large Twitter diffusion datasets, VBHP enjoys linear fitting time with the data size and fast convergence rate, and provides more accurate predictions than those of state-of-the-art approaches. The novel ELBO is also demonstrated to exceed the common one in model selection.

Comparison with the Approach in Chapter 3. The approach in Chapter 3 similarly exploits the Hawkes process branching structure and the finite support assumption of the triggering kernel, while it builds on Gibbs sampling. It considers only high-probability triggering relationships in computations for acceleration. However, those relationships are updated in each iteration, which is less efficient than our precomputing them. Besides, our variational inference schema enjoys faster convergence than that of the Gibbs sampling based method.

Both methods in this chapter and the last chapter are designed based on classical approximate Bayesian inference, we propose a new approximate inference technique which outperforms classical ones in the next chapter.

# Quantile Propagation for Wasserstein-Approximate Gaussian Processes 

Approximate inference techniques are the cornerstone of probabilistic methods based on Gaussian process (GP) priors. Despite this, most work approximately optimizes standard divergence measures such as the Kullback-Leibler (KL) divergence, which lack the basic desiderata for the task at hand, while chiefly offering merely technical convenience. In this chapter, we develop an efficient approximate Bayesian scheme that minimizes a specific class of Wasserstein distances (WDs), which we refer to as the $L_{2}$ WD. Our method overcomes some of the shortcomings of the KL divergence for approximate inference with GP models. The contents of this chapter is organized as below:
(i) In Section 5.1, we first introduce the background and motivation of the new work.
(ii) In Section 5.2, we develop quantile propagation (QP), an approximate inference algorithm for models with GP priors and factorized likelihoods. Like the expectation propagation (EP) algorithm, QP does not directly minimize global distances between high-dimensional distributions. Instead, QP estimates a fully coupled Gaussian posterior by iteratively minimizing local divergences between two particular marginal distributions. As these marginals are univariate, QP boils down to an iterative quantile function matching procedure (rather than moment matching as in EP) - hence we term our inference scheme quantile propagation. We derive the updates for the approximate likelihood terms and show that while the QP mean estimates match those of EP, the variance estimates are lower for QP.
(iii) In Section 5.3, we show that like EP, QP satisfies the locality property, meaning that it is sufficient to employ univariate approximate likelihood terms, and that the updates can thereby be performed efficiently using only the marginal distributions. Consequently, although our method employs a more complex divergence than that of EP ( $L_{2} \mathrm{WD}$ vs KL), it has the same computational complexity, after the precomputation of certain (data independent) lookup
tables.
(iv) In Section 5.4, we employ eight real-world datasets and compare our method to EP and variational Bayes (VB) on the tasks of binary classification and Poisson regression. We find that in terms of predictive accuracy, QP performs similarly to EP but is superior to VB. In terms of predictive uncertainty, however, we find QP superior to both EP and VB, thereby supporting our claim that QP alleviates variance over-estimation associated with the KL divergence when approximating short-tailed distributions [Minka, 2005; Jylänki et al., 2011; Heess et al., 2013].

### 5.1 Introduction

Gaussian process (GP) models have attracted the attention of the machine learning community due to their flexibility and their capacity to measure uncertainty. They have been widely applied to learning tasks such as regression [Matheron, 1963], classification [Williams and Barber, 1998; Hensman et al., 2015] and stochastic point process modeling [Møller et al., 1998; Zhang et al., 2019]. However, exact Bayesian inference for GP models is intractable for all but the Gaussian likelihood function. To address this issue, various approximate Bayesian inference methods have been proposed, such as Markov Chain Monte Carlo [MCMC, see e.g. Neal, 1997], the Laplace approximation [Williams and Barber, 1998], variational inference [Jordan et al., 1999; Opper and Archambeau, 2009] and expectation propagation [Opper and Winther, 2000; Minka, 2001c].

The existing approach most relevant to this work is expectation propagation (EP), which approximates each non-Gaussian likelihood term with a Gaussian by iteratively minimizing a set of local forward Kullback-Leibler (KL) divergences. As shown by Gelman et al. [2017], EP can scale to very large datasets. However, EP is not guaranteed to converge, and is known to over-estimate posterior variances [Minka, 2005; Jylänki et al., 2011; Heess et al., 2013] when approximating a short-tailed distribution. By over-estimation, we mean that the approximate variances are larger than the true variances so that more distribution mass lies in the ineffective domain. Interestingly, many popular likelihoods for GPs results in short-tailed posterior distributions, such as Heaviside and probit likelihoods for GP classification and Laplacian, Student's t and Poisson likelihoods for GP regression.

The tendency to over-estimate posterior variances is an inherent drawback of the forward KL divergence for approximate Bayesian inference. More generally, several authors have pointed out that the KL divergence can yield undesirable results such as (but not limited to) over-dispersed or under-dispersed posteriors [Dieng et al., 2017; Li and Turner, 2016; Hensman et al., 2014].

As an alternative to the KL, optimal transport metrics-such as the Wasserstein distance [WD, Villani, 2008, §6]-have seen a recent boost of attention. The WD is a natural distance between two distributions, and has been successfully employed in tasks such as image retrieval [Rubner et al., 2000], text classification [Huang et al., 2016] and image fusion [Courty et al., 2016]. Recent work has begun to employ the

WD for inference, as in Wasserstein generative adversarial networks [Arjovsky et al., 2017], Wasserstein variational inference [Ambrogioni et al., 2018] and Wasserstein auto-encoders [Tolstikhin et al., 2017]. In contrast to the KL divergence, the WD is computationally challenging [Cuturi, 2013], especially in high dimensions [Bonneel et al., 2015], in spite of its intuitive formulation and excellent performance.

### 5.2 Quantile Propagation

We now propose our new approximation algorithm which, as summarized in Algorithm 2 (Appendix), employs an $L_{2}$ WD based projection rather than the forward KL divergence projection of EP. Although QP employs a more complex divergence, it has the same computational complexity as EP, with the following caveat. To match the speed of EP, it is necessary to precompute sets of (data independent) lookup tables. Once precomputed, the resulting updates are only a constant factor slower than EP - a modest price to pay for optimizing a divergence which is challenging even to evaluate. Appendix C. 10 provides further details on the precomputation and use of these tables.

As stated in Proposition 1, minimizing $\mathrm{W}_{2}^{2}\left(\widetilde{q}\left(f_{i}\right), \mathcal{N}\left(f_{i}\right)\right)$ is equivalent to minimizing the $L_{2}$ distance between quantile functions of $\widetilde{q}\left(f_{i}\right)$ and $\mathcal{N}\left(f_{i}\right)$, so we refer to our method as quantile propagation (QP). This section focuses on deriving local updates for the site functions and analyzing their relationships with those of EP. Later in Section 5.3 , we show the locality property of QP, meaning that the site function $t(f)$ has a univariate parameterization and so the local update can be efficiently performed using marginals only.

### 5.2.1 Convexity of $L_{p}$ Wasserstein Distance

We first show $\mathrm{W}_{p}^{p}\left(\widetilde{q}(f), \mathcal{N}\left(f \mid \mu, \sigma^{2}\right)\right)$ to be strictly convex in $\mu$ and $\sigma$. Formally, we have:

Theorem 2. Given two probability measures in $\mathcal{M}_{+}^{1}(\mathbb{R})$ : a Gaussian $\mathcal{N}\left(\mu, \sigma^{2}\right)$ with mean $\mu$ and standard deviation $\sigma>0$, and an arbitrary measure $\widetilde{q}, W_{p}^{p}(\widetilde{q}, \mathcal{N})$ is strictly convex in $\mu$ and $\sigma$.

Proof. See Appendix C.4.

### 5.2.2 Minimization of $L_{2}$ WD

An advantage of using the $L_{p}$ WD with $p=2$, rather than other choices of $p$, is the computational efficiency it admits in the local updates. As we show in this section, optimizing the $L_{2}$ WD yields neat analytical updates of the optimal $\mu^{\star}$ and $\sigma^{\star}$ that require only univariate integration and the CDF of $\widetilde{q}(f)$. In contrast, other $L_{p}$ WDs lack convenient analytical expressions. Nonetheless, other $L_{p}$ WDs may have useful properties, the study of which we leave to future work.

The optimal parameters $\mu^{\star}$ and $\sigma^{\star}$ corresponding to the $L_{2} \mathrm{WD} \mathrm{W}_{2}^{2}\left(\widetilde{q}, \mathcal{N}\left(\mu, \sigma^{2}\right)\right)$ can be obtained using Proposition 1. Specifically, we employ the quantile function reformulation of $\mathrm{W}_{2}^{2}\left(\widetilde{q}, \mathcal{N}\left(\mu, \sigma^{2}\right)\right)$, and zero its derivatives w.r.t. $\mu$ and $\sigma$. The results provided below are derived in section C.1:

$$
\begin{align*}
& \mu^{\star}=\mu_{\widetilde{q}} \\
& \sigma^{\star}=\sqrt{2} \int_{0}^{1} F_{\widetilde{q}}^{-1}(y) \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y=1 / \sqrt{2 \pi} \int_{-\infty}^{\infty} e^{-\left[\operatorname{erf}^{-1}\left(2 F_{\widetilde{q}}(f)-1\right)\right]^{2}} \mathrm{~d} f \tag{5.1}
\end{align*}
$$

Interestingly, the update for $\mu$ matches that of EP, namely the expectation under $\widetilde{q}$. However, for the standard deviation we have the difficulty of deriving the CDF $F_{\widetilde{q}}$. If a closed form expression is available, we can apply numerical integration to compute the optimal standard deviation; otherwise, we may use sampling based methods to approximate it. In our experiments we employ the former.

### 5.2.3 Properties of the Variance Update

Given the update equations in the previous section, here we show that the standard deviation estimate of QP , denoted as $\sigma_{\mathrm{QP}}$, is less or equal to that of EP , denoted as $\sigma_{\mathrm{EP}}$, when projecting the same tilted distribution to the Gaussian space.

Theorem 3. The variances of the Gaussian approximation to a univariate tilted distribution $\widetilde{q}(f)$ as estimated by $Q P$ and $E P$ satisfy $\sigma_{Q P}^{2} \leq \sigma_{E P}^{2}$.

Proof. See Appendix C.5.
Corollary 3.1. The variances of the site functions updated by EP and $Q P$ satisfy: $\widetilde{\sigma}_{Q P}^{2} \leq \widetilde{\sigma}_{E P}^{2}$ and the variances of the approximate posterior marginals satisfy $\sigma_{q, Q P}^{2} \leq \sigma_{q, E P}^{2}$.

Proof. Since the cavity distribution is unchanged, we can calculate the variance of the site function as per Equation (2.11) and conclude that the variance of the site function also satisfies $\widetilde{\sigma}_{\mathrm{QP}}^{2} \leq \widetilde{\sigma}_{\mathrm{EP}}^{2}$. Moreover as per the definition of the cavity distribution in Section 2.2.5, the approximate marginal distribution is proportional to the product of the cavity distribution and the site function $q\left(f_{i}\right) \propto q^{\backslash i}\left(f_{i}\right) t\left(f_{i}\right)$, which are two Gaussian distributions. By the product of Gaussians formula (Equation (2.11)), we know the variance of $q\left(f_{i}\right)$ estimated by EP as $\sigma_{q, \mathrm{EP}}^{2}=\left(\widetilde{\sigma}_{\mathrm{EP}}^{-2}+\sigma_{\backslash i}^{-2}\right)^{-1}=\sigma_{\mathrm{EP}}^{2}$ and similarly $\sigma_{q, \mathrm{QP}}^{2}=\sigma_{\mathrm{QP}}^{2}$, where $\sigma_{\mathrm{EP}}^{2}$ and $\sigma_{\mathrm{QP}}^{2}$ are defined in Theorem 3. Thus, there is $\sigma_{q, \mathrm{QP}}^{2} \leq \sigma_{q, \mathrm{EP}}^{2}$.

Corollary 3.2. The predictive variances of latent functions at $\boldsymbol{x}_{*}$ by EP and QP satisfy: $\sigma_{Q P}^{2}\left(f\left(\boldsymbol{x}_{*}\right)\right) \leq \sigma_{E P}^{2}\left(f\left(\boldsymbol{x}_{*}\right)\right)$.

Proof. The predictive variance of the latent function was analyzed in [Rasmussen and Williams, 2005, Equation (3.61)]: $\sigma^{2}\left(f_{*}\right)=k_{*}-\boldsymbol{k}_{*}^{\mathrm{T}}(K+\widetilde{\Sigma})^{-1} \boldsymbol{k}_{*}$,
where we define $f_{*}=f\left(\boldsymbol{x}_{*}\right)$ and $k_{*}=k\left(\boldsymbol{x}_{*}, \boldsymbol{x}_{*}\right)$, and let $\boldsymbol{k}_{*}=\left(k\left(\boldsymbol{x}_{*}, \boldsymbol{x}_{i}\right)\right)_{i=1}^{N}$ be the (column) covariance vector between the test data $\boldsymbol{x}_{*}$ and the training data $\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{N}$.

After updating parameters of the site function $t_{i}\left(f_{i}\right)$, the predictive variance can be written as (details in Appendix C.9):

$$
\sigma_{\text {new }}^{2}\left(f_{*}\right)=k_{*}-\boldsymbol{k}_{*}^{\mathrm{T}} A \boldsymbol{k}_{*}+\boldsymbol{k}_{*}^{\mathrm{T}} \boldsymbol{s}_{i} \boldsymbol{s}_{i}^{\mathrm{T}} \boldsymbol{k}_{*} /\left[\left[\widetilde{\sigma}_{i, \text { new }}^{2}-\widetilde{\sigma}_{i}^{2}\right)^{-1}+A_{i i}\right],
$$

where $\widetilde{\sigma}_{i \text {,new }}^{2}$ is the site variance updated by EP or QP, $A=(K+\widetilde{\Sigma})^{-1}$ and $s_{i}$ is the $i^{\prime}$ s column of $A$. Since $\widetilde{\sigma}_{i, \mathrm{QP}}^{2} \leq \widetilde{\sigma}_{i, \mathrm{EP}}^{2}$, we have $\sigma_{\mathrm{QP}}^{2}\left(f_{*}\right) \leq \sigma_{\mathrm{EP}}^{2}\left(f_{*}\right)$.

Remark. We compared variance estimates of EP and QP assuming the same cavity distribution. Proving analogous statements for the fixed points of the EP and QP algorithms is more challenging, however, and we leave this to future work, while providing empirical support for these analogous statements in Figure 5.1(a) and Figure 5.1(b).

### 5.3 Locality Property

In this section we detail the central result on which our QP algorithm is based upon, which we refer to as the locality property. That is, the optimal site function $t_{i}$ is defined only in terms of the single corresponding latent variable $f_{i}$, and thereby and similarly to EP, it admits a simple and efficient sequential update of each individual site approximation.

### 5.3.1 Review: Locality Property of EP

We provide a brief review of the locality property of EP for GP models; for more details see Seeger [2005]. We begin by defining the general site function $t_{i}(\boldsymbol{f})$ in terms of all of the latent variables, and the cavity and the tilted distributions as $q^{\backslash i}(\boldsymbol{f}) \propto p(\boldsymbol{f}) \prod_{j \neq i} \widetilde{t}_{j}(\boldsymbol{f})$ and $\widetilde{q}(\boldsymbol{f}) \propto q^{i}(\boldsymbol{f}) p\left(y_{i} \mid f_{i}\right)$, respectively. To update $t_{i}(\boldsymbol{f})$, EP matches a multivariate Gaussian distribution $\mathcal{N}(f)$ to $\widetilde{q}(f)$ by minimizing the KL divergence $\operatorname{KL}(\widetilde{q} \| \mathcal{N})$, which is further rewritten as (see details in Appendix C.6.1):

$$
\begin{equation*}
\operatorname{KL}(\widetilde{q} \| \mathcal{N})=\operatorname{KL}\left(\widetilde{q}_{i} \| \mathcal{N}_{i}\right)+\mathbb{E}_{\widetilde{q}_{i}}\left[\operatorname{KL}\left(q_{\backslash i \mid i}^{\backslash i} \| \mathcal{N}_{\backslash i \mid i}\right)\right], \tag{5.2}
\end{equation*}
$$

where and hereinafter, $\backslash i \mid i$ denotes the conditional distribution of $f_{\backslash i}$ (taking $f_{i}$ out of f) given $f_{i}$, namely, $q_{\backslash i \mid i}^{\backslash i}=q^{\backslash i}\left(f_{\backslash i} \mid f_{i}\right)$ and $\mathcal{N} \backslash i \mid i=\mathcal{N}\left(f_{\backslash i} \mid f_{i}\right)$. Note that $q_{\backslash i \mid i}^{\backslash i}$ and $\mathcal{N}_{i \mid i}$ in the second term in Equation (5.2) are both Gaussian, and so setting them equal to one another causes that term to vanish. Furthermore, it is well known that the term $\operatorname{KL}\left(\widetilde{q}_{i} \| \mathcal{N}_{i}\right)$ is minimized w.r.t. the parameters of $\mathcal{N}_{i}$ by matching the first and second moments of $\widetilde{q}_{i}$ and $\mathcal{N}_{i}$. Finally, according to the usual EP logic, we recover the site function $t_{i}(f)$ by dividing the optimal Gaussian $\mathcal{N}(f)$ by the cavity $q^{i}(f)$ :

$$
\begin{equation*}
t_{i}(f) \propto \mathcal{N}(f) / q^{\backslash i}(f)=\underline{\mathcal{N}}\left(\left.\boldsymbol{f}_{i}\right|_{i}\right) \mathcal{N}\left(f_{i}\right) /\left(q^{\backslash i}\left(f_{i} \mid f_{i}\right) q^{i}\left(f_{i}\right)\right)=\mathcal{N}\left(f_{i}\right) / q^{\backslash i}\left(f_{i}\right) . \tag{5.3}
\end{equation*}
$$

Here we can see the optimal site function $t_{i}\left(f_{i}\right)$ relies solely on the local latent variable $f_{i}$, so it is sufficient to assume a univariate expression for site functions. Besides, the
site function can be efficiently updated by using the marginals $\widetilde{q}\left(f_{i}\right)$ and $\mathcal{N}\left(f_{i}\right)$ only, namely, $\left.t_{i}\left(f_{i}\right) \propto\left(\min _{\mathcal{N}_{i}} \operatorname{KL}\left(\widetilde{q}_{i} \| \mathcal{N}_{i}\right)\right) / q^{i}\left(f_{i}\right)\right)$.

### 5.3.2 Locality Property of QP

This section proves the locality property of QP, which turns out to be rather more involved to show than is the case for EP. We first prove the following theorem, and then follow the same procedure as for EP (Equation (5.3)).

Theorem 4. Minimization of $W_{2}^{2}(\widetilde{q}(f), \mathcal{N}(f))$ w.r.t. $\mathcal{N}(f)$ results in $q^{\backslash i}\left(f_{\backslash} \mid f_{i}\right)=\mathcal{N}\left(f_{\backslash i} \mid f_{i}\right)$.
Proof. See Appendix C.6.
Theorem 5 (Locality Property of QP). For GP models with factorized likelihoods, QP requires only univariate site functions, and so yields efficient updates using only marginal distributions.

Proof. We apply the same steps as in Equation (5.3) for the EP case to QP and we conclude that the site function $t_{i}\left(f_{i}\right) \propto \mathcal{N}\left(f_{i}\right) / q^{\backslash i}\left(f_{i}\right)$ relies solely on the local latent variable $f_{i}$. And as per Equation (C.15) (Appendix C.6), $\mathcal{N}\left(f_{i}\right)$ is estimated by $\min _{\mathcal{N}_{i}} W_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)$, so the local update only uses marginals and can perform efficiently.

Benefits of the Locality Property. The locality property admits an analytically economic form for the site function $t_{i}\left(f_{i}\right)$, requiring a parameterization that depends on a single latent variable. In addition, this also yields a significant reduction in the computational complexity, as only marginals are involved in each local update. In contrast, if QP (or EP) had no such a locality property, estimating the mean and the variance would involve integrals w.r.t. high-dimensional distributions, with a significantly higher computational cost should closed form expressions be unavailable.

### 5.4 Experiments

In this section, we compare the QP, EP and variational Bayes [VB, Opper and Archambeau, 2009] algorithms for binary classification and Poisson regression. The experiments employ eight real world datasets and aim to compare relative accuracy of the three methods, rather than optimizing the absolute performance. The implementations of EP and VB in Python are publicly available [GPy, since 2012], and our implementation of QP is based on that of EP. For both EP and QP, we stop local updates, i.e., , the inner loop in Algorithm 2 (Appendix), when the root mean squared change in parameters is less than $10^{-6}$. In the outer loop, the GP hyper-parameters are optimized by L-BFGS-B [Byrd et al., 1995] with a maximum of $10^{3}$ iterations and a relative tolerance of $10^{-9}$ for the function value. VB is also optimized by L-BFGS-B with the same configuration. Parameters shared by the three methods are initialized to be the same.

### 5.4.1 Binary Classification

Benchmark Data. We perform binary classification experiments on the five real world datasets employed by Kuss and Rasmussen [2005]: Ionosphere (IonoS), Wisconsin Breast Cancer, Sonar [Dua and Graff, 2017], Leptograpsus Crabs and Pima Indians Diabetes [Ripley, 1996]. We use two additional UCI datasets as further evidence: Glass and Wine [Dua and Graff, 2017]. As the Wine dataset has three classes, we conduct binary classification experiments on all pairs of classes. We summarize the dataset size and data dimensions in Table 5.1.

Table 5.1: Results on benchmark datasets. The first three columns give dataset names, the number of instances $m$ and the number of features $n$. The table records the test errors (TEs) and the negative test log-likelihoods (NTLLs). The top section is on the benchmark datasets employed by Kuss and Rasmussen [2005] and the middle section uses additional datasets. The bottom section shows Poisson regression results. * indicates that QP outperforms EP in more than $90 \%$ of experiments consistently.

| Data | $\mathrm{TE}\left(\times 10^{-2}\right)$ |  |  | NTLL $\left(\times 10^{-3}\right)$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | EP | QP | VB | EP | QP | VB |
| IonoS 35134 | $7.9 \pm \pm .5$ | $7.9{ }_{ \pm 0.5}$ | $18.9 \pm 6.9$ | $215.9 \pm 8.4$ | $215.9 \pm 8.5$ | $337.4_{ \pm 70.8}$ |
| Cancer 6839 | $3.2 \pm 0.2$ | $3.2 \pm 0.2$ | $3.1{ }_{ \pm 0.2}$ | $88.2 \pm 3.1$ | $88.2{ }_{ \pm 3.1}^{*}$ | $88.9 \pm 19.1$ |
| Pima 7327 | $20.3{ }_{ \pm 1.0}$ | $20.3{ }_{ \pm 1.0}$ | $21.9 \pm 0.4$ | $424.7_{ \pm 13.0}$ | $424.0{ }_{ \pm 13.2}^{*}$ | $450.3_{ \pm 2.6}$ |
| Crabs 2007 | $2.7{ }_{ \pm 0.5}$ | $2.7{ }_{ \pm 0}$ | $3.7 \pm 0$ | $64.4 \pm$ | $64.3 \pm 8.4$ | $164.7_{ \pm 7.5}$ |
| Sonar 20860 | $14.0{ }_{ \pm 1.1}$ | $14.0 \pm 1.1$ | $25.7 \pm 3.9$ | $306.7_{ \pm 10.8}$ | 306.2* ${ }_{\text {+10.9 }}$ | $693.1_{ \pm 0.0}$ |
| lass 21410 | $1.1_{ \pm 0.4}$ | $1.0 \pm \pm 0.4$ | $2.6 \pm 0$. | $29.5 \pm 5.4$ | 29.0 ${ }_{ \pm 5.5}^{*}$ | $79.5 \pm$ ¢6.3 |
| Wine1 13013 | $1.5_{ \pm 0.5}$ | $1.5{ }_{ \pm 0.5}$ | $1.7 \pm 0.6$ | $48.0 \pm 3.4$ | 47.4 ${ }_{ \pm 3.4}^{*}$ | $83.9{ }_{ \pm 5.2}$ |
| Wine2 10713 | 0.0 ${ }_{ \pm 0.0}$ | 0.0 ${ }_{ \pm 0.0}$ | $0.0 \pm 0.0$ | $18.0_{ \pm 1.2}$ | $17.8{ }_{ \pm 1.2}^{*}$ | $26.7{ }_{ \pm 1.9}$ |
| Wine3 11913 | $2.0_{ \pm 1.0}$ | $2.0_{ \pm 1.0}$ | $1.2_{ \pm 0.7}$ | $52.1_{ \pm 5.6}$ | $51.8{ }_{ \pm 5.6}^{*}$ | $69.4 \pm 5.0$ |

Mining $1121118.6_{ \pm 27.0} 118.6_{ \pm 27.0} 170.3_{ \pm 15.9} 1606.8_{ \pm 116.3} \mathbf{1 6 0 6 . 5}_{ \pm 116.3} 2007.3_{ \pm 119.8}$
Note: Wine1: Class 1 vs. 2. Wine2: Class 1 vs. 3. Wine3: Class 2 vs. 3.
Prediction. We predict the test labels using models optimized by EP, QP and VB on the training data. For a test input $\boldsymbol{x}_{*}$ with a binary target $y_{*}$, the approximate predictive distribution is written as: $q\left(y_{*} \mid \boldsymbol{x}_{*}\right)=\int_{-\infty}^{\infty} p\left(y_{*} \mid f_{*}\right) q\left(f_{*}\right) \mathrm{d} f_{*}$ where $f_{*}=f\left(\boldsymbol{x}_{*}\right)$ is the value of the latent function at $x_{*}$. We use the probit likelihood for the binary classification task, which admits an analytical expression for the predictive distribution and results in a short-tailed posterior distribution. Correspondingly, the predicted label $\hat{y}_{*}$ is determined by thresholding the predictive probability at $1 / 2$.

Performance Evaluation. To evaluate the performance, we employ two measures: the test error (TE) and the negative test log-likelihood (NTLL). The TE and the NTLL quantify the prediction accuracy and uncertainty, respectively. Specifically, they are defined as $\left(\sum_{i=1}^{m}\left|y_{*, i}-\hat{y}_{*, i}\right| / 2\right) / m$ and $-\left(\sum_{i=1}^{m} \log q\left(y_{*, i} \mid \boldsymbol{x}_{*, i}\right)\right) / m$, respectively, for a set of test inputs $\left\{x_{*, i}\right\}_{i=1}^{m}$, test labels $\left\{y_{*, i}\right\}_{i=1}^{m}$, and the predicted labels $\left\{\hat{y}_{*, i}\right\}_{i=1}^{m}$. Lower values indicate better performance for both measures.

Experiment Settings. In the experiments, we randomly split each dataset into 10 folds, each time using 1 fold for testing and the other 9 folds for training, with features standardized to zero mean and unit standard deviation. We repeat this 100 times for a random seed ranging 0 through 99. As a result, there are a total of 1,000 experiments for each dataset. We report the average and the standard deviation of the above metrics over the 100 rounds.

Results. The evaluation results are summarized in Table 5.1. The top section presents the results on the datasets employed by Kuss and Rasmussen [2005], whose reported TEs match ours as expected. While QP and EP exhibit similar TEs on these datasets, QP is superior to EP in terms of the NTLL. VB under-performs both EP and QP on all datasets except Cancer. The middle section of Table 5.1 shows the results on additional datasets. The TEs are again similar for EP and QP, while QP has lower NTLLs. Again, VB performs worst among the three methods. To emphasize the difference between NTLLs of EP and QP, we mark with an asterisk those results in which QP outperforms EP in more than $90 \%$ of the experiments. Furthermore, we visualize the predictive variances of QP in comparison with those of EP in Figure 5.1(a)., which shows that the variances of QP are always less than or equal to those of EP, thereby providing empirical evidence of QP alleviating the over-estimation of predictive variances associated with the EP algorithm.

### 5.4.2 Poisson Regression

Data and Settings. We perform a Poisson regression experiment to further evaluate the performance of our method. The experiment employs the coal-mining disaster dataset [Jarrett, 1979] which has 190 data points indicating the time of fatal coal mining accidents in the United Kingdom from 1851 to 1962. To generate training and test sequences, we randomly assign every point of the original sequence to either a training or test sequence with equal probability, and this is repeated 200 times (random seeds $0, \cdots, 199$ ), resulting in 200 pairs of training and test sequences. We use the TE and the NTLL to evaluate the performance of the model on the test dataset. The NTLL has the same expression as that of the Binary classification experiment, but with a different predictive distribution $q\left(y_{*} \mid \boldsymbol{x}_{*}\right)$. The TE is defined slightly differently as $\left(\sum_{i=1}^{m}\left|y_{*, i}-\hat{y}_{*, i}\right|\right) / m$. To make the rate parameter of the Poisson likelihood nonnegative, we use the square link function [Flaxman et al., 2017; Walder and Bishop, 2017], and as a result, the likelihood becomes $p\left(y \mid f^{2}\right)$. We use this link function because it is more mathematically convenient than the exponential function: the EP and QP update formulas, and the predictive distribution $q\left(y_{*} \mid \boldsymbol{x}_{*}\right)$ are available in Appendices C.3.2 and C.8, respectively.

Results. The means and the standard deviations of the evaluation results are reported in the last row of Table 5.1. Compared with EP, QP yields lower NTLL, which implies a better fitting performance of QP to the test sequences. We also provide the predictive variances in Figure 5.1(b)., in the variance of QP is once again seen to be less than or equal to that of EP. This experiment further supports our claim that QP alleviates the problem with EP of over-estimation of the predictive variance.


Figure 5.1: A scatter plot of the predictive variances of latent functions on test data, for EP and QP. The diagonal dash line represents equivalence. We see that the predictive variance of QP is always less than or equal to that of EP.

Finally, once again we find that both EP and QP outperform VB.

### 5.5 Conclusions

We have proposed QP as the first efficient $L_{2}-\mathrm{WD}$ based approximate Bayesian inference method for Gaussian process models with factorized likelihoods. Algorithmically, QP is similar to EP but uses the $L_{2}$ WD instead of the forward KL divergence for estimation of the site functions. When the likelihood factors are approximated by a Gaussian form we show that QP matches quantile functions rather than moments as in EP. Furthermore, we show that QP has the same mean update but a smaller variance than that of EP, which in turn alleviates the over-estimation by EP of the posterior variance in practice. Crucially, QP has the same favorable locality property as EP, and thereby admits efficient updates. Our experiments on binary classification and Poisson regression have shown that QP can outperform both EP and variational Bayes. Approximate inference with WD is promising but hard to compute, especially for continuous multivariate distributions. We believe our work paves the way for further practical approaches to WD-based inference.

Limitations and Future Work. Although we have presented properties and advantages of our method, it is still worth pointing out its limitations. First, our method does not provide a methodology for hyper-parameter optimization that is consistent with our proposed WD minimization framework. Instead, for this purpose, we rely on optimization of EP's marginal likelihood. We believe this is one of the reasons for the small performance differences between QP and EP. Furthermore, the computational efficiency of our method comes at the price of additional memory requirements and the look-up tables may exhibit instabilities on high-dimensional data. To overcome
these limitations, future work will explore alternatives to hyper-parameter optimization, improvements on numerical computation under the current approach and a variety of WD distances under a similar algorithm framework.

Applications to Bayesian Hawkes Processes. Applying QP to Bayesian Hawkes processes is a challenging task. This goal first requires solving a less difficult problem: using the EP or QP algorithm for Bayesian Poisson processes. We didn't explore QP for Poisson processes but for Poisson regression. Even for EP, the application to Poisson processes is missing. The main reason is that Poisson processes involve generally analytically intractable integrals in the likelihood and therefore are hard to handle. Further to this problem, applying QP to Bayesian Hawkes processes will face new problems, such as a high computational burden.

In the next chapter, we study a simple and robust frequentist estimation framework beyond the Bayesian field. The framework builds on the kernel maximum moment restriction and is applicable to a wide range of models.

# Kernel Maximum Moment Restriction for Instrumental Variable Regression 

In this chapter, we propose a simple estimation framework for conditional moment restriction (CMR) models. We focus on instrumental variable (IV) regression models and design the method based on the kernelized maximum moment restriction (MMR). We organize the contents as below:
(i) In Section 6.1, we introduce the background of IV regression and present an overview of our proposed method.
(ii) In Section 6.2, we elaborate the prerequisites of our method: the detailed settings of IV regression, generalized method of moment which is widely applied to IV regression estimation and related to our method, and the kernelized MMR framework which is formulated by maximizing the interaction between the residual and the instruments belonging to a unit ball in a reproducing kernel Hilbert space (RKHS).
(iii) In Section 6.3, we propose the simple framework for IV regression estimation based on the kernelized MMR. Different from two-step optimization in most of existing methods, our method reformulates the IV regression estimation as a single-step empirical risk minimization problem, where the risk depends on the reproducing kernel on the instrument and can be estimated by a U-statistic or V-statistic. We then present two practical algorithms by considering two modern machine learning models, neural networks and kernel machines, in the framework.
(iv) In Section 6.4, we present an efficient hyper-parameter selection procedure.
(v) In Section 6.5, we analyze consistency and asymptotic normality of our estimator in both parametric and non-parametric settings.
(vi) In Section 6.6, we demonstrate the advantages of our framework over existing ones using experiments on both synthetic and real-world data.

### 6.1 Introduction

Instrumental variables (IV) have become standard tools for economists, epidemiologists, and social scientists to uncover causal relationships from observational data [Angrist and Pischke, 2008; Klungel et al., 2015]. Randomization of treatments or policies has been perceived as the gold standard for such tasks, but is generally prohibitive in many real-world scenarios due to time constraints or ethical concerns. When treatment assignment is not randomized, it is generally impossible to discern between the causal effect of treatments and spurious correlations that are induced by unobserved factors. Instead, the use of IV enables the investigators to incorporate natural variation through an IV that is associated with the treatments, but not with the outcome variable, other than through its effect on the treatments. In economics, for instance, the season-of-birth was used as an IV to study the return from schooling, which measures causal effect of education on labor market earning [Card, 1999]. In genetic epidemiology, the idea to use genetic variants as IVs, known as Mendelian randomization, has also gained increasing popularity [Burgess et al., 2017b, 2020].

To overcome these drawbacks, we propose a simple framework which views nonlinear IV regression as an empirical risk minimization (ERM) problem with Ustatistic or V-statistic. This framework is based on (i) the maximum moment restriction (MMR) by Muandet et al. [2020a], which develops an equivalent form of moment conditions by a reproducing kernel Hilbert space (RKHS), and (ii) a U/V-statistics approximation technique [Serfling, 1980] for the form. We call our framework MMRIV. Based on this formulation, we can solve the nonlinear IV problem by a single-step optimization with an empirical risk.

Our MMR-IV framework has the following advantages:
(i) A closed-form solution of the optimization problem is available. Also, for neural networks, we can simply apply common algorithms such as SGD.
(ii) We can interpret MMR-IV as an analogy for Gaussian processes, and consequently, we develop an efficient hyper-parameter selection procedure.
(iii) We prove consistency and asymptotic normality of estimators by MMR-IV in both parametric and non-parametric settings.

All of these advantages come from the empirical risk minimization form of MMR-IV. To the best of our knowledge, we do not find any other method that can achieve all of these. Further, our experiments show MMR-IV has better performance with synthetic data, and it provides appropriate interpretation of real data.

### 6.2 Preliminaries

Generalized method of moment (GMM). The aforementioned conditions imply that $\mathbb{E}[\varepsilon \mid Z]=0$ for $P_{Z}$-almost all $z$. This is a conditional moment restriction (CMR) which we can use to estimate $f$ [Newey, 1993]. For any measurable function $h$, the CMR implies a continuum of unconditional moment restrictions [Lewis and Syrgkanis, 2018;

Bennett et al., 2019] ${ }^{1}$ :

$$
\begin{equation*}
\mathbb{E}[(Y-f(X)) h(Z)]=0 \tag{6.1}
\end{equation*}
$$

That is, there exists an infinite number of moment conditions, each of which is indexed by the function $h$. As a result, learning with Equation (6.1) is challenging. Although the asymptotic efficiency of the estimator can in principle improve when we consider increasingly many moment conditions, it was observed that the excessive number of moments can be harmful in practice [Andersen and Sorensen, 1996] because its finite-sample bias increases with the number of moment conditions [Newey and Smith, 2004]. Hence, traditional works in econometrics often select a finite number of moment conditions for estimation based on the generalized method of moments (GMM) [Hansen, 1982; Hall et al., 2005]. Unfortunately, an adhoc choice of moments can potentially lead to a loss of efficiency or even a loss of identification [Dominguez and Lobato, 2004]. For this reason, subsequent works advocate an incorporation of all moment restrictions simultaneously in different ways such as the method of sieves [de Jong, 1996; Donald et al., 2003] and a continuum of moment restrictions [Carrasco and Florens, 2000; Carrasco et al., 2007; Carrasco, 2012; Carrasco and Florens, 2014], among others.

One of the key questions in econometrics is which moment conditions should be used as a basis for estimating the function $f$ [Donald and Newey, 2001; Hall, 2005]. In this work, we show that, for the purpose of estimating $f$, it is sufficient to restrict $h$ to be within a unit ball of a RKHS of real-valued functions on $Z$.

### 6.2.1 Maximum Moment Restriction

Throughout this paper, we assume that $h$ is a real-valued function on $\mathcal{Z}$ which belongs to a RKHS $\mathcal{H}_{k}$ endowed with a reproducing kernel $k: \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$. The RKHS $\mathcal{H}_{k}$ satisfies two important properties: (i) for all $z \in \mathcal{Z}$ and $h \in \mathcal{H}_{k}$, we have $k(z, \cdot) \in \mathcal{H}_{k}$ and (ii) (reproducing property of $\left.\mathcal{H}_{k}\right) h(z)=\langle h, k(z, \cdot)\rangle_{\mathcal{H}_{k}}$ where $k(z, \cdot)$ is a function of the second argument. Furthermore, we define $\Phi_{k}(z)$ as a canonical feature map of $z$ in $\mathcal{H}_{k}$. It follows from the reproducing property that $k\left(z, z^{\prime}\right)=\left\langle\Phi_{k}(z), \Phi_{k}\left(z^{\prime}\right)\right\rangle_{\mathcal{H}_{k}}$ for any $z, z^{\prime} \in \mathcal{Z}$, i.e., an inner product between the feature maps of $z$ and $z^{\prime}$ can be evaluated implicitly through the kernel evaluation. Every positive definite kernel $k$ uniquely determines the RKHS for which $k$ is a reproducing kernel [Aronszajn, 1950]. For detailed exposition on kernel methods, see, e.g., Schölkopf and Smola [2002], Berlinet and Thomas-Agnan [2004], and Muandet et al. [2017].

Instead of (6.1), we define a risk in terms of a maximum moment restriction (MMR) [Muandet et al., 2020a]:

$$
\begin{equation*}
R_{k}(f):=\sup _{h \in \mathcal{H}_{k},\|h\| \leq 1}(\mathbb{E}[(Y-f(X)) h(Z)])^{2} . \tag{6.2}
\end{equation*}
$$

That is, instead of considering all measurable functions $h$ as instruments, we only restrict to functions that lie within a unit ball of the RKHS $\mathcal{H}_{k}$. The risk is then defined

[^2]as a maximum value of the moment restriction with respect to this function class. The benefits of our formulation are two-fold. First, it is computationally intractable to learn $f$ from (6.1) using all measurable functions as instruments. By restricting the function class to a unit ball of the RKHS, the problem becomes computationally tractable, as will be shown in Lemma 1 below. Second, this restriction still preserves the consistency of parameter estimated using $R_{k}(f)$. In some sense, the RKHS is a sufficient class of instruments for the nonlinear IV problem (cf. Theorem 6).

A crucial insight for our approach is that the population risk $R_{k}(f)$ has an analytic solution.

Lemma 1 (Muandet et al. [2020a], Thm 3.3). Assume that $\mathbb{E}\left[(Y-f(X))^{2} k(Z, Z)\right]<\infty$. Then, we have

$$
\begin{equation*}
R_{k}(f)=\mathbb{E}\left[(Y-f(X))\left(Y^{\prime}-f\left(X^{\prime}\right)\right) k\left(Z, Z^{\prime}\right)\right] \tag{6.3}
\end{equation*}
$$

where $\left(X^{\prime}, Y^{\prime}, Z^{\prime}\right)$ is an independent copy of $(X, Y, Z)$.
We assume throughout that the reproducing kernel $k$ is integrally strictly positive definite (ISPD).
Assumption 1. The kernel $k$ is continuous, bounded (i.e., $\sup _{z \in \mathcal{Z}} \sqrt{k(z, z)}<\infty$ ) and satisfies the condition of integrally strictly positive definite (ISPD) kernels, i.e., for any function $g$ that satisfies $0<\|g\|_{2}^{2}<\infty$, we have $\iint_{\mathcal{Z}} g(z) k\left(z, z^{\prime}\right) g\left(z^{\prime}\right) \mathrm{d} z \mathrm{~d} z^{\prime}>0$.

The assumption is also related to the notion of characteristic and universal kernels; see, e.g., Simon-Gabriel and Schölkopf [2018]. More details on ISPD kernels are given in Appendix D.1. We further assume the identification for the minimizer of $R_{k}(f)$.
Assumption 2. Consider the function space $\mathcal{F}$ and $f^{*} \in \operatorname{argmin}_{f \in \mathcal{F}} R_{k}(f)$. Then for any $g \in \mathcal{F}$ with $|\mathbb{E}[g(X)]|<\infty, \mathbb{E}\left[g(X)-f^{*}(X) \mid Z\right]=0$ implies $g=f^{*}$.

A sufficient condition for identification follows from the completeness property of $X$ for Z [D'Haultfoeuille, 2011], e.g., the conditional distribution of $X$ given $Z$ belongs to the exponential family [Newey and Powell, 2003]. Provided identification, it is straightforward to obtain consistency.

### 6.3 Our Method

We propose to learn $f$ by minimizing $R_{k}(f)$ in (6.3). To this end, we define an optimal function $f^{*}$ as a minimizer of the above population risk with respect to a function class $\mathcal{F}$ of real-valued functions on $\mathcal{X}$, i.e.,

$$
f^{*} \in \operatorname{argmin}_{f \in \mathcal{F}} R_{k}(f) .
$$

It is instructive to note that population risk $R_{k}$ depends on the choice of the kernel $k$. Based on Assumption 1 and Lemma 1, we obtain the following result, which is a special case of Muandet et al. [2020a, Theorem 3.2], showing that $R_{k}(f)=0$ if and only if $f$ satisfies the original conditional moment restriction (see Appendix D.2.2 for the proof).

Theorem 6. Assume that the kernel $k$ is integrally strictly positive definite (ISPD). Then, for any real-valued measurable function $f, R_{k}(f)=0$ if and only if $\mathbb{E}[Y-f(X) \mid z]=0$ for $P_{Z}$-almost all $z$.

Theorem 6 holds as long as the kernel $k$ belongs to a class of ISPD kernels. Hence, it allows for more flexibility in terms of the kernel choice. Moreover, it is not difficult to show that $R_{k}(f)$ is strictly convex in $f$ (see Appendix D.2.3).

### 6.3.1 Empirical Risk Minimization with U, V-Statistics

The previous results pave the way for an empirical risk minimization (ERM) framework [Vapnik, 1998] to be used in our work. That is, given an i.i.d. sample $\left\{\left(x_{i}, y_{i}, z_{i}\right)\right\}_{i=1}^{n} \sim$ $P^{n}(X, Y, Z)$ of size $n$, an empirical estimate of the risk $R_{k}(f)$ can be obtained as

$$
\begin{equation*}
\widehat{R}_{U}(f):=\sum_{1 \leq i \neq j \leq n}^{n} \frac{\left(y_{i}-f\left(x_{i}\right)\right)\left(y_{j}-f\left(x_{j}\right)\right) k\left(z_{i}, z_{j}\right)}{n(n-1)}, \tag{6.4}
\end{equation*}
$$

which is in the form of U-statistic [Serfling, 1980, Section 5]. Alternatively, an empirical risk based on V-statistic can also be used, i.e.,

$$
\begin{equation*}
\widehat{R}_{V}(f):=\sum_{i, j=1}^{n} \frac{\left(y_{i}-f\left(x_{i}\right)\right)\left(y_{j}-f\left(x_{j}\right)\right) k\left(z_{i}, z_{j}\right)}{n^{2}} . \tag{6.5}
\end{equation*}
$$

Both forms of empirical risk can be used as a basis for a consistent estimation of $f$. The advantage of (6.4) is that it is a minimum-variance unbiased estimator with appealing asymptotic properties, whereas (6.5) is a biased estimator of the population risk (6.3), i.e., $\mathbb{E}\left[\widehat{R}_{V}\right] \neq R_{k}$. However, the estimator based on V-statistics employs a full sample and hence may yield better estimate of the risk than the U-statistic counterpart.

Let $x:=\left[x_{1}, \ldots, x_{n}\right]^{\top}, y:=\left[y_{1}, \ldots, y_{n}\right]^{\top}$ and $z:=\left[z_{1}, \ldots, z_{n}\right]^{\top}$ be column vectors. Let $K_{\mathbf{z}}$ be the kernel matrix $K(\boldsymbol{z}, \boldsymbol{z})=\left[k\left(z_{i}, z_{j}\right)\right]_{i j}$ evaluated on the instruments $\mathbf{z}$. Then, both (6.4) and (6.5) can be rewritten as

$$
\begin{equation*}
\widehat{R}_{V(U)}(f)=(\boldsymbol{y}-f(\boldsymbol{x}))^{\top} W_{V(U)}(\boldsymbol{y}-f(\boldsymbol{x})), \tag{6.6}
\end{equation*}
$$

where $f(x):=\left[f\left(x_{1}\right), \ldots, f\left(x_{n}\right)\right]^{\top}$ and $W_{V(U)} \in \mathbb{R}^{n \times n}$ is a symmetric weight matrix that depends on the kernel matrix $K_{z}$. Specifically, the weight matrix $W$ corresponding to (6.4) is given by $W_{U}=\left(K_{z}-\operatorname{diag}\left(k\left(z_{1}, z_{1}\right), \ldots, k\left(z_{n}, z_{n}\right)\right)\right) /(n(n-1))$ where $\operatorname{diag}\left(a_{1}, \ldots, a_{n}\right)$ denotes an $n \times n$ diagonal matrix whose diagonal elements are $a_{1}, \ldots, a_{n}$. As shown in Appendix D.2.5, $W_{U}$ is indefinite and may cause problematic inferences. The weight matrix $W$ for (6.5) is given by $W_{V}:=K_{z} / n^{2}$ which is a positive definite matrix for the ISPD kernel $k$. Finally, our objective (6.6) also resembles the well-known generalized least regression with correlated noise [Kariya and Kurata, 2004, Chapter 2] where the covariance matrix is the Z-dependent invertible matrix $W_{V(U)}^{-1}$.

Based on (6.4) and (6.5), we estimate $f^{*}$ by minimizing the regularized empirical risk over a function class $\mathcal{F}$ :

$$
\hat{f}_{V(U)} \in \operatorname{argmin}_{f \in \mathcal{F}} \widehat{R}_{V(U)}(f)+\lambda \Omega(f)
$$

where $\lambda>0$ is a regularization constant satisfying $\lim _{n \rightarrow \infty} \lambda=0$, and $\Omega(f)$ is the regularizer. Since $\hat{f}_{U}$ and $\hat{f}_{V}$ minimizes objectives which are regularized U-statistic and V-statistic, they can be viewed a specific form of M-estimators; see, e.g., Van der Vaart [2000, Ch. 5]. In this work, we focus on the V-statistic empirical risk and provide practical algorithms when $\mathcal{F}$ is parametrized by deep neural networks (NNs) and an RKHS of real-valued functions.

Kernelized GMM. We may view the objective (6.6) from the GMM perspective [Hall, 2005]. The assumption that the instruments $Z$ are exogenous implies that $\mathbb{E}\left[\Phi_{k}(Z) \varepsilon\right]=0$ where $\Phi_{k}$ denotes the canonical feature map associated with the kernel $k$. This gives us an infinite number of moments, $g(f)=\Phi_{k}(Z)(Y-$ $f(X))$. Hence, we can write the sample moments as $\hat{g}(f)=(1 / n) \sum_{i=1}^{n} \Phi_{k}\left(z_{i}\right)\left(y_{i}-\right.$ $f\left(x_{i}\right)$ ). The intuition behind GMM is to choose a function $f$ that sets these moment conditions as close to zero as possible, motivating the objective function $J(f):=$ $\|\hat{g}(f)\|_{\mathcal{H}_{k}}^{2}=\langle\hat{g}(f), \hat{g}(f)\rangle_{\mathcal{H}_{k}}=\frac{1}{n^{2}} \sum_{i, j=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right)\left\langle\Phi_{k}\left(z_{i}\right), \Phi_{k}\left(z_{j}\right)\right\rangle_{\mathcal{H}_{k}}\left(y_{j}-f\left(x_{j}\right)\right)=$ $(y-f(x))^{\top} W_{V}(y-f(x))$. Hence, our objective (6.6) defined using V-statistic is a special case of the GMM objective when the weighting matrix is the identity operator. Carrasco et al. [2007, Ch. 6] shows that the optimal weighting operator is given in terms of the inversed covariance operator.

### 6.3.2 Practical MMR-IV Algorithms

A workflow of our algorithm based on $\widehat{R}_{V}$ is summarized in Algorithm 1; we leave the $\widehat{R}_{U}$ based method to future work to solve the inference issues caused by indefinite $W_{U}$. We provide examples of the class $\mathcal{F}$ in both parametric and non-parametric settings below.

Deep neural networks. In the parametric setting, the function class $\mathcal{F}$ can often be expressed as $\mathcal{F}_{\Theta}=\left\{f_{\theta}: \theta \in \Theta\right\}$ where $\Theta \subseteq \mathbb{R}^{m}$ denotes a parameter space. We consider a very common nonlinear model in machine learning $f(x)=W_{0} \Phi(x)+b_{0}$ where $\Phi: x \longmapsto \sigma_{h}\left(W_{h} \sigma_{h-1}\left(\cdots \sigma_{1}\left(W_{1} x\right)\right)\right)$ denotes a nonlinear feature map of a depth- $h$ NN. Here, $W_{i}$ for $i=1, \ldots, h$ are parameter matrices and each $\sigma_{i}$ denotes the entry-wise activation function of the $i$-th layer. In this case, $\theta=\left(b_{0}, W_{0}, W_{1}, \ldots, W_{h}\right)$. As a result, we can rewrite $\hat{f}_{V}$ in terms of their parameters as

$$
\hat{\theta}_{V} \in \arg \min _{\theta \in \Theta} \widehat{R}_{V}\left(f_{\theta}\right)+\lambda\|\theta\|_{2}^{2}
$$

where $f_{\theta} \in \mathcal{F}_{\Theta}$. We denote $\theta^{*} \in \arg \min _{\theta \in \Theta} R_{k}\left(f_{\theta}\right)$. In what follows, we refer to this algorithm as MMR-IV (NN); see Algorithm 4 in Appendix D.5.1.

Kernel machines. In a non-parametric setting, the function class $\mathcal{F}$ becomes an infinite dimensional space. In this work, we consider $\mathcal{F}$ to be an RKHS $\mathcal{H}_{l}$ of real-valued functions on $\mathcal{X}$ with a reproducing kernel $l: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. Then, the regularized solution can be obtained by $\arg \min _{f \in \mathcal{H}_{l}} \widehat{R}_{V}(f)+\lambda\|f\|_{\mathcal{H}_{l}}^{2}$. As per the representer theorem, any optimal $\hat{f}$ admits a form $\hat{f}(x)=\sum_{i=1}^{n} \alpha_{i} l\left(x, x_{i}\right)$ for some $\left(\alpha_{1}, \ldots, \alpha_{n}\right) \in \mathbb{R}^{n}$ [Schölkopf et al., 2001b], and based on this representation, we

```
Algorithm 1 MMR-IV
Input: Dataset \(D=\{x, y, z\}\), kernel function \(k\) with parameters \(\theta_{k}\), a function class
    \(\mathcal{F}\), regularization functional \(\Omega(\cdot)\), and regularization parameter \(\lambda\)
Output: The estimate of \(f^{*}\) in \(\mathcal{F}\).
    Compute the kernel matrix \(K=k\left(z, z ; \theta_{k}\right)\).
    Define the residual \(\varepsilon(f)=y-f(x)\).
    \(\hat{f}_{\lambda} \leftarrow \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n^{2}} \varepsilon(f)^{\top} \operatorname{K\varepsilon }(f)+\lambda \Omega\left(\|f\|_{\mathcal{F}}\right)\)
    return \(\hat{f}_{\lambda}\)
```

rewrite the objective as

$$
\begin{equation*}
\hat{f}_{V}=\arg \min _{\boldsymbol{\alpha} \in \mathbb{R}^{n}}(\boldsymbol{y}-L \boldsymbol{\alpha})^{\top} W_{V}(\boldsymbol{y}-L \boldsymbol{\alpha})+\lambda \boldsymbol{\alpha}^{\top} L \boldsymbol{\alpha}, \tag{6.7}
\end{equation*}
$$

where $L=\left[l\left(x_{i}, x_{j}\right)\right]_{i j}$ is the kernel matrix on $x$. For U-statistic version, the quadratic program (6.7) substitutes indefinite $W_{U}$ for $W_{V}$, so it may not be positive definite. The value of $\lambda$ needs to be sufficiently large to ensure that (6.7) is definite. On the other hand, the V-statistic based estimate (6.7) is definite for all non-zero $\lambda$ since $W_{V}$ is positive semi-definite. Thus, the optimal $\widehat{\alpha}$ can be obtained by solving the first-order stationary condition and if $L$ is positive definite, the solution has a closed form expression, $\widehat{\boldsymbol{\alpha}}=\left(L W_{V} L+\lambda L\right)^{-1} L W_{V} y$. Thus, we will focus on the V-statistic version in our experiments. In the following, we refer to this algorithm as MMR-IV (RKHS).

Nyström approximation. The MMR-IV (RKHS) algorithm is computationally costly for large datasets as it requires a matrix inversion. To improve the scalability, we resort to Nyström approximation [Williams and Seeger, 2001] to accelerate the matrix inversion in $\widehat{\boldsymbol{\alpha}}=\left(L W_{V} L+\lambda L\right)^{-1} L W_{V} y$. First, we randomly select a subset of $m(\ll n)$ samples from the original dataset and construct the corresponding submatrices of $W_{V}$, namely, $W_{m m}$ and $W_{n m}$ based on this subset. Second, let $V$ and $U$ be the eigenvalue vector and the eigenvector matrix of $W_{m m}$. Then, the Nyström approximation is obtained as $W_{V} \approx \widetilde{U} \widetilde{V} \widetilde{U}^{\top}$ where $\widetilde{U}:=\sqrt{\frac{m}{n}} W_{n m} U V^{-1}$ and $\widetilde{V}:=\frac{n}{m} V$. We finally apply the Woodbury formula [Flannery et al., 1992, p. 75] to obtain

$$
\begin{aligned}
\left(L W_{V} L+\lambda L\right)^{-1} L W_{V} & =L^{-1}\left(W_{V}+\lambda L^{-1}\right)^{-1} W_{V} \\
& \approx \lambda^{-1}\left[I-\widetilde{U}\left(\lambda^{-1} \widetilde{U}^{\top} L \widetilde{U}+\widetilde{V}^{-1}\right)^{-1} \widetilde{U}^{\top} \lambda^{-1} L\right] \widetilde{U} \widetilde{V} \widetilde{U}^{\top}
\end{aligned}
$$

We will refer to this algorithm as MMR-IV (Nyström); Algorithm 3 in Appendix D.5.1. The runtime complexity of this algorithm is $O\left(n m^{2}+n^{2}\right)$.

### 6.4 Hyper-parameter Selection

We develop a convenient hyperparameter selection method for MMR-IV (Nyström) with our objective (6.7) and the V-statistic. This method benefits from the fact that our ERM form can be interpreted as a stochastic model with a Gaussian process, inspired
by Vehtari et al. [2016]. This approach is an analog of parameter selection in the ordinary kernel regression. However, we need to apply additional analysis due to the weight matrix $W_{V}$ in (6.7) with V-statistics. Full details are provided in Appendix D.3.

Gaussian process (GP) interpretation. We start the GP interpretation by defining the energy-based likelihood $p(D \mid f(x)) \propto \exp \left(-\hat{R}_{V}(f) / 2\right) \propto N\left(f(x) \mid \boldsymbol{y}, K_{z}^{-1}\right)$, where $D:=\{x, y, z\}$ denotes the dataset. Then we assign a GP prior to $f$, i.e., $f(x) \sim$ $\mathrm{GP}(0, \delta l(x, x))$ with a real constant $\delta>0$,s and the posterior distribution of $f(x)$ is straight-forwardly derived based on the Gaussian-like prior and likelihood,

$$
p(f(\boldsymbol{x}) \mid D)=N(f(\boldsymbol{x}) \mid \boldsymbol{c}, C) \propto p(f(\boldsymbol{x})) p(D \mid f(\boldsymbol{x})),
$$

where $\boldsymbol{c}=C K_{z} y$ and $C=\left(K_{z}+(\delta L)^{-1}\right)^{-1} \approx \delta\left[L-L \widetilde{U}\left(n^{2} \delta \widetilde{U}^{\top} L \widetilde{U}+\widetilde{V}^{-1}\right)^{-1} \widetilde{U}\left(n^{2} \delta\right) L\right]$ by Nyström approximation. The connections between the GP model and the regularized $\widehat{R}_{V}(f)$ are elaborated in Appendix D. 3 and summarized as Theorem 7. In a word, maximization of $p(f(x) \mid D)$ is alternative to minimization of (6.7) and the Bayesian inference is a substitute for the frequentist prediction.

Theorem 7. Given $\delta=\left(\lambda n^{2}\right)^{-1}$ and $\hat{f}=\operatorname{argmin}_{f \in \mathcal{H}_{l}}$ Equation (6.7), there are
(i) $\operatorname{argmax}_{f(x)} p(f(x) \mid D)=\hat{f}(x)$;
(ii) prediction at $\boldsymbol{x}^{*}$ : $\operatorname{argmax}_{f\left(x_{*}\right)} p\left(f\left(\boldsymbol{x}_{*}\right) \mid D\right)=\hat{f}\left(\boldsymbol{x}_{*}\right)$.

Analytical cross-validation error. Now we derive the analytical error of the leave- $M$-out cross validation (LMOCV) from the perspective of the GP. We split the whole dataset $D$ into training and development datasets, denoted as $D_{t r}$ and $D_{d e}:=\left\{\boldsymbol{x}_{d e}, y_{d e}, z_{d e}\right\}$ respectively, where $D_{d e}$ has $M$ triplets of data points. Given $D_{t r}$, the predictive probability on $D_{d e}$ can be obtained by Bayes' rules $p\left(f\left(\boldsymbol{x}_{d e}\right) \mid D_{t r}\right) \propto$ $\frac{p\left(f\left(x_{d e}\right) \mid D\right)}{p\left(D_{d e} f f\left(x_{d e}\right)\right)}$, where $p\left(f\left(x_{d e}\right) \mid D\right)=N\left(c_{d e}, C_{d e}\right)$ with $c_{d e}, C_{d e}$ the mean and covariance of $f\left(\boldsymbol{x}_{d e}\right)$ in $p(f(\boldsymbol{x}) \mid D)$ and $p\left(D_{d e} \mid f\left(\boldsymbol{x}_{d e}\right)\right) \propto N\left(\boldsymbol{y}_{d e}, K_{d e}^{-1}\right), K_{d e}=k\left(\boldsymbol{z}_{d e}, \boldsymbol{z}_{d e}\right)$. By this result, the function estimated on the training set can be represented w.r.t. that on the whole dataset. It turns out that $p\left(f\left(\boldsymbol{x}_{d e}\right) \mid D_{t r}\right)=N(\boldsymbol{b}, B)$ where $B^{-1}=C_{d e}^{-1}-K_{d e}$ and $\boldsymbol{b}=B\left(C_{d e}^{-1} \boldsymbol{c}_{d e}-K_{d e} \boldsymbol{y}_{d e}\right)$, and by (ii) in Theorem 7, the error of $m$ repeated LMOCV is

$$
\begin{equation*}
\text { LMOCV Error : }=\sum_{i=1}^{m}\left(r^{(i)}\right)^{\top} K_{d e}^{(i)} r^{(i)}, \tag{6.8}
\end{equation*}
$$

where (i) denotes the experiment index and the residual $r^{(i)}:=\boldsymbol{b}^{(i)}-\boldsymbol{y}_{M}^{(i)}=(I-$ $\left.C_{d e}^{(i)} K_{d e}^{(i)}\right)^{-1}\left(c_{d e}^{(i)}-y_{d e}^{(i)}\right)$. The analytical error enables an efficient parameter selection with the CV procedure.

### 6.5 Consistency and Asymptotic Normality

We provide the consistency and asymptotic normality of $\widehat{f}_{V}$. We also develop the same results for $\widehat{f}_{U}$, but we defer them to the appendix due to space limitation. All proofs are presented in the appendix.

### 6.5.1 Consistency

We first show the consistency of $\hat{f}_{V}$, which depends on the uniform convergence of the risk functions. The result holds regardless of the shape of $\Omega(f)$, so we can utilize the regularization $\|\theta\|_{2}^{2}$ which is common for NN but non-convex in terms of $f$.
Theorem 8 (Consistency of $\hat{f}_{V}$ ). Assume that $\mathbb{E}\left[|Y|^{2}\right]<\infty, \mathbb{E}\left[\sup _{f \in \mathcal{F}}|f(X)|^{2}\right]<\infty, \mathcal{F}$ is compact, Assumption 1, 2 hold, $\Omega(f)$ is a bounded function and $\lambda \xrightarrow{\mathrm{p}} 0$. Then $\hat{f}_{V} \xrightarrow{\mathrm{p}} f^{*}$.

If the $\Omega(f)$ is convex in $f$, the consistency can be obtained more easily by Newey and McFadden [1994, Theorem 2.7]. In this case, we can avoid several conditions. We provide an additional result with this setting in Appendix D.2.6.

### 6.5.2 Asymptotic Normality

We analyze asymptotic normality of the estimator $\hat{f}_{V}$, which is important to advanced statistical analysis such as tests. Here, we investigate two different cases: the estimator has finite- and infinite-dimension.

Finite-dimension case. We consider the $\hat{f}_{V}$ is characterized by a finite-dimensional parameter from a parameter space $\Theta$. We rewrite the regularized V-statistic risk as a compact form $\widehat{R}_{V, \lambda}\left(f_{\theta}\right):=\frac{1}{n^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n} h_{\theta}\left(u_{i}, u_{j}\right)+\lambda \Omega(\theta), h_{\theta}\left(u_{i}, u_{j}\right):=\left(y_{i}-\right.$ $\left.f\left(x_{i}\right)\right)\left(y_{j}-f\left(x_{j}\right)\right) k\left(z_{i}, z_{j}\right)$, and consider $R_{k}\left(f_{\theta}\right)$ is uniquely minimized at $\theta^{*} \in \Theta$.
Theorem 9 (Asymptotic normality of $\hat{\theta}_{V}$ ). Suppose that $f_{\theta}$ and $\Omega(\theta)$ are twice continuously differentiable about $\theta$, $\Theta$ is compact, $H=\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]$ is non-singular, $\mathbb{E}\left[|Y|^{2}\right]<$ $\infty, \mathbb{E}\left[\sup _{\theta \in \Theta}\left|f_{\theta}(X)\right|^{2}\right]<\infty, \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}^{2}\right]<\infty, \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} f_{\theta}(X)\right\|_{F}^{2}\right]<$ $\infty, \sqrt{n} \lambda \xrightarrow{\mathrm{P}} 0, R_{k}\left(f_{\theta}\right)$ is uniquely minimized at $\theta^{*}$ which is an interior point of $\Theta$, and Assumption 1 holds. Then, $\sqrt{n}\left(\hat{\theta}_{V}-\theta^{*}\right) \rightsquigarrow N\left(\mathbf{0}, \Sigma_{V}\right)$ holds, where

$$
\Sigma_{V}=4 H^{-1} \operatorname{diag}\left(\mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right) H^{-1}
$$

and $\rightsquigarrow$ denotes a convergence in law.
Infinite-dimension case. We show the asymptotic normality of an infinite-dimensional estimator $\hat{f}_{V}$. That is, we show that an error of $\hat{f}_{V}$ weakly converges to a Gaussian process that takes values in a function space $\mathcal{H}_{l}$. We set $\Omega(f)=\|f\|_{\mathcal{H}_{l}}^{2}$ and consider a minimizer: $f_{\lambda_{0}}^{*} \in \operatorname{argmin}_{f \in \mathcal{F}} R_{k}(f)+\lambda_{0}\|f\|_{\mathcal{H}_{l}}^{2}$ with arbitrary $\lambda_{0}>0$. We also define $\mathcal{N}(\varepsilon, \mathcal{H},\|\cdot\|)$ as an $\varepsilon$-covering number of a set $\mathcal{H}$ in terms of $\|\cdot\|$. Then, we obtain the following:

Theorem 10. Suppose Assumption 1 holds, $l$ is a bounded kernel, $k$ is a uniformly bounded function, and $\lambda-\lambda_{0}=o\left(n^{-1 / 2}\right)$ holds. Also, suppose that $\mathcal{X}, \mathcal{Z}$, and $\mathcal{Y}$ are compact spaces, and there exists $s \in(0,2)$ and a constant $C_{H}>0$ such that $\log \mathcal{N}\left(\varepsilon, \mathcal{H}_{l},\|\cdot\|_{L^{\infty}}\right) \leq C_{H} \varepsilon^{-s}$ for any $\varepsilon \in(0,1)$. Then, there exists a Gaussian process $\mathbb{G}_{P}^{*}$ such that $\sqrt{n}\left(\hat{f}_{V}-f_{\lambda_{0}}^{*}\right) \rightsquigarrow$ $\mathrm{G}_{P}^{*}$ in $\mathcal{H}_{l}$.

This result allows statistical inference on functional estimators such as kernel machines. Although there are many conditions, all of them are valid for many well-known kernels; see Appendix D.2.11 for examples.

### 6.6 Experimental Results

We present the experimental results in a wide range of settings for IV estimation. Following Lewis and Syrgkanis [2018] and Bennett et al. [2019], we consider both low and high-dimensional scenarios. In our experiments, we compare our algorithms to the following baselines: (i) DirectNN (ii) 2SLS (iii) Poly2SLS (iv) DeepIV [Hartford et al., 2017] (v) KernelIV [Singh et al., 2019] (vi) GMM+NN (vii) AGMM [Lewis and Syrgkanis, 2018] (viii) DeepGMM [Bennett et al., 2019]. (ix) AGMM-K [Dikkala et al., 2020]. We refer readers to Appendix D. 5 for more experimental details.

### 6.6.1 Low-dimensional Scenarios

Following Bennett et al. [2019], we employ the following data generation process:

$$
Y=f^{*}(X)+e+\delta, \quad X=Z_{1}+e+\gamma,
$$

where $Z:=\left(Z_{1}, Z_{2}\right) \sim \operatorname{Uniform}\left([-3,3]^{2}\right), e \sim \mathcal{N}(0,1)$, and $\gamma, \delta \sim \mathcal{N}\left(0,0.1^{2}\right)$. In words, $Z$ is a two-dimensional IV, but only the first instrument $Z_{1}$ has an effect on X . The variable $e$ is the confounding variable that creates the correlation between $X$ and the residual $Y-f^{*}(X)$. We vary the true function $f^{*}$ between the following cases to enrich the datasets: (i) sin: $f^{*}(x)=\sin (x)$. (ii) step: $f^{*}(x)=\mathbb{1}_{\{x \geq 0\}}$. (iii) abs: $f^{*}(x)=|x|$. (iv) linear: $f^{*}(x)=x$. We consider both small-sample ( $n=200$ ) and large-sample ( $n=2000$ ) regimes.

Table 6.1 reports the results for the large-sample regime (and Table D. 1 in Appendix D. 5 for the small-sample regime). Our findings are as follows: (i) MMR-IVs perform reasonably well in most cases. (ii) The linear methods (2SLS and Poly2SLS) perform best when the linearity assumption is satisfied. (iii) Some nonlinear but com-

Table 6.1: The mean square error (MSE) $\pm$ one standard deviation in the large-sample regime ( $n=2000$ ).

| Algorithm | abs | True Function $f^{*}$ |  |  |
| :--- | :---: | :---: | :---: | :---: |
|  | linear | sin | step |  |
| DirectNN | $.116 \pm .000$ | $.035 \pm .000$ | $.189 \pm .000$ | $.199 \pm .000$ |
| 2SLS | $.522 \pm .000$ | $.000 \pm .000$ | $.254 \pm .000$ | $.050 \pm .000$ |
| Poly2SLS | $.083 \pm .000$ | $.000 \pm .000$ | $.133 \pm .000$ | $.039 \pm .000$ |
| GMM+NN | $.318 \pm .000$ | $.044 \pm .000$ | $.694 \pm .000$ | $.500 \pm .000$ |
| AGMM | $.600 \pm .001$ | $.025 \pm .000$ | $.274 \pm .000$ | $.047 \pm .000$ |
| DeepIV | $.247 \pm .004$ | $.056 \pm .003$ | $.165 \pm .003$ | $.038 \pm .001$ |
| DeepGMM | $.027 \pm .009$ | $.005 \pm .001$ | $.160 \pm .025$ | $.025 \pm .006$ |
| KernelIV | $.019 \pm .000$ | $.009 \pm .000$ | $.046 \pm .000$ | $.026 \pm .000$ |
| AGMM-K | $181 \pm .000$ | $2.34 \pm .000$ | $19.4 \pm .000$ | $4.13 \pm .000$ |
| MMR-IV (NN) | $.011 \pm .002$ | $.005 \pm .000$ | $.153 \pm .019$ | $.040 \pm .004$ |
| MMR-IV (Nys) | $.011 \pm .001$ | $.001 \pm .000$ | $.006 \pm .002$ | $.020 \pm .002$ |



Figure 6.1: Runtime comparison in the large-sample regime ( $n=2000$ ). The computational time of parameter selection is excluded from the comparison. AGMM-K (No Nyström) and MMR-IV (RKHS) overlap due to the same runtime.
plicated methods (GMM+NN, DeepIV and DeepGMM) are unstable due to the sensitivity to hyper-parameters. (iv) We suspect that the AGMM-K performance is unsatisfactory because the hyper-parameter selection is not flexible enough. On the other hand, MMR-IV (Nyström) has the advantage of adaptive hyper-parameter selection (cf. Section 6.4). Appendix D.5.2 provides more details.

We also record the runtimes of all methods on the large-sample regime and report them in Figure 6.1. Compared to the NN-based methods, i.e., AGMM, DeepIV, DeepGMM, our MMR-IV (NN) is the most computationally efficient method, which is clearly a result of a simpler objective. Using a minimax optimization between two NNs, AGMM is the least efficient method. DeepGMM and DeepIV are more efficient than AGMM, but are less efficient than MMR-IV (NN). Lastly, all three RKHS-based methods, namely, KernelIV, AGMM-K and MMR-IV (RKHS), have similar computational time. All three methods are observed to scale poorly on large datasets.

### 6.6.2 High-dimensional Structured Scenarios

In high-dimensional setting, we employ the same data generating process as in Section 6.6.1. We consider only the absolute function for $f^{*}$, but map $Z, X$, or both $X$ and $Z$ to MNIST images (784-dim) [LeCun et al., 1998]. Let us denote the original outputs in Section 6.6 .1 by $X^{\text {low }}, Z^{\text {low }}$ and let $\pi(u):=\operatorname{round}(\min (\max (1.5 u+5,0), 9))$ be a transformation function mapping inputs to an integer between 0 and 9 , and let $\mathrm{RI}(d)$ be a function that selects a random MNIST image from the digit class $d$. We set $n=10,000$. Then, the scenarios we consider are
(i) MNIST $_{\mathrm{Z}}: Z \leftarrow \operatorname{RI}\left(\pi\left(Z_{1}^{\text {low }}\right)\right)$,
(ii) MNIST $_{\mathbf{X}}: X \leftarrow \operatorname{RI}\left(\pi\left(X^{\text {low }}\right)\right)$,



Number of IVs
Value of $c_{1}$
Value of $c_{2}$


Figure 6.2: The MSE of different methods on Mendelian randomization experiments as we vary the numbers of instruments (left), the strength of confounders to exposures $c_{1}$ (middle), and the strength of confounders to instruments $c_{2}$ (right). The MSE is obtained from 10 repetitions of the experiment.
(iii) MNIST MZ $: X \leftarrow \operatorname{RI}\left(\pi\left(X^{\text {low }}\right)\right), Z \leftarrow \operatorname{RI}\left(\pi\left(Z_{1}^{\text {low }}\right)\right)$.

We report the results in Table 6.2. We summarize our findings: (i) MMR-IV (NN) performs well across scenarios. (ii) MMR-IV (Nyström) fails except MNIST ${ }_{Z}$, because the ARD kernel with PCA are not representative enough for MNIST . . (iii) DeepIV $^{\text {P }}$ does not work when $X$ is high-dimensional, similar to Bennett et al. [2019]. (iv) The two-step methods (Poly2SLS and Ridge2SLS) has large errors because the first-stage regression from $Z$ to $X$ is ill-posed. Appendix D.5.3 provides more details.

### 6.6.3 Mendelian Randomization

We demonstrate our method in the setting of Mendelian randomization which relies on genetic variants that satisfy the IV assumptions. The "exposure" $X$ and outcome $Y$ are univariate and generated from the simulation process by Hartwig et al. [2017]:

$$
Y=\beta X+c_{1} e+\delta, \quad X=\sum_{i=1}^{m} \alpha_{i} Z_{i}+c_{2} e+\gamma
$$

where $Z \in \mathbb{R}^{d^{\prime}}$ with each entry $Z_{i} \sim B\left(2, p_{i}\right), p_{i} \sim \operatorname{unif}(0.1,0.9), e \sim \mathcal{N}(0,1)$, $\alpha_{i} \sim \operatorname{unif}\left(\left[0.8 / d^{\prime}, 1.2 / d^{\prime}\right]\right)$, and $\gamma, \delta \sim \mathcal{N}\left(0,0.1^{2}\right) . Z_{i} \sim B\left(2, p_{i}\right)$ mimics the frequency of an individual getting one or more genetic variants. The parameters $\beta, c_{1}$ control the strength of exposures and confounders to outcomes, while $c_{2}, \alpha_{i}$ control the strength of instruments and confounders to exposures. We set $\alpha_{i} \sim \operatorname{unif}\left(\left[0.8 / d^{\prime}, 1.2 / d^{\prime}\right]\right)$ so that as the number of IVs increases, each instrument becomes weaker while the overall strength of instruments ( $\sum_{i}^{d^{\prime}} \alpha_{i}$ ) remains constant.

In Mendelian randomization, it is known that genetic variants may act as weak IVs [Kuang et al., 2020; Hartford et al., 2020; Burgess et al., 2020], so this experiment aims to evaluate the sensitivity of different methods to the number of instruments ( $d^{\prime}$ ) and confounder strengths $\left(c_{1}, c_{2}\right)$. We consider three experiments: (i) $d^{\prime}=8,16,32$; (ii) $c_{1}=0.5,1,2$; (iii) $c_{2}=0.5,1,2$; unmentioned parameters use default values: $\beta=1$,

Table 6.2: The mean square error (MSE) $\pm$ one standard deviation on high-dimensional structured data. We run each method 10 times.

| Algorithm | MNIST $_{\mathrm{z}}$ | Setting <br> $\mathrm{MNIST}_{\mathrm{x}}$ | MNIST $_{\mathrm{xz}}$ |
| :--- | :---: | :---: | :---: |
| DirectNN | $.134 \pm .000$ | $.229 \pm .000$ | $.196 \pm .011$ |
| 2SLS | $.563 \pm .001$ | $>1000$ | $>1000$ |
| Ridge2SLS | $.567 \pm .000$ | $.431 \pm .000$ | $.705 \pm .000$ |
| GMM+NN | $.121 \pm .004$ | $.235 \pm .002$ | $.240 \pm .016$ |
| AGMM | $.017 \pm .007$ | $.732 \pm .107$ | $.529 \pm .163$ |
| DeepIV | $.114 \pm .005$ | $\mathrm{n} / \mathrm{a}$ | $\mathrm{n} / \mathrm{a}$ |
| DeepGMM | $.038 \pm .004$ | $.315 \pm .130$ | $.333 \pm .168$ |
| AGMM-K+NN | $.021 \pm .007$ | $1.05 \pm .366$ | $.327 \pm .192$ |
| MMR-IV (NN) | $.024 \pm .006$ | $.124 \pm .021$ | $.130 \pm .009$ |
| MMR-IV (Nys) | $.015 \pm .002$ | $.442 \pm .000$ | $.425 \pm .002$ |

$d^{\prime}=16, c_{1}=1, c_{2}=1$. We draw 10,000 samples for the training, validation and test sets, respectively, and train MMR-IV (Nyström) only on the training set. Other settings are the same as those of the low-dimensional scenario.

Figure 6.2 depicts the experimental results. Overall, 2SLS performs well on all settings due to the linearity assumption, except particular sensitivity to the number of (weak) instruments, which is a well-known property of 2SLS [Angrist and Pischke, 2008]. Although imposing no such assumption, MMR-IVs perform competitively and even more stably, since the information of instruments is effectively captured by the kernel $k$ and we only need to deal with a simple objective, and also the analytical CV error plays an essential role. Section D.5.5 contains additional findings.

### 6.6.4 Application on the Vitamin D data

Lastly, we apply our algorithm to the Vitamin D data [Sjolander and Martinussen, 2019, Sec. 5.1]. The data were collected from a 10-year study on 2571 individuals aged 40-71 and 4 variables are employed: age (at baseline), filaggrin (binary indicator of filaggrin mutations), VitD (Vitamin D level at baseline) and death (binary indicator of death during study). The goal is to evaluate the potential effect of VitD on death. We follow Sjolander and Martinussen [2019] by controlling age in the analyses, using filaggrin as instrument, and then applying the MMR-IV (Nyström) algorithm. Sjolander and Martinussen [2019] modeled the effect of VitD on death by a generalized linear model and found the effect is insignificant by 2SLS ( $p$-value on the estimated coefficient is 0.13 with the threshold of 0.05). More details can be found in Appendix D.5.6. The estimated effect is illustrated in Figure D. 1 in the appendix. We observe that: (i) By using instruments, both our method and Sjolander and Martinussen [2019] output more reasonable results compared with those without instruments: a low VitD level at a young age has a slight effect on death, but a more adverse effect at an old age [Meehan and Penckofer, 2014]; (ii) Unlike Sjolander and Martinussen [2019], our
method allows more flexible non-linearity for causal effect.

### 6.7 Conclusion

Learning causal relations when hidden confounders are present is a cornerstone of decision making. IV regression is a standard tool to tackle this task, but currently faces challenges in nonlinear settings. The present work presents a simple framework that overcomes some of these challenges. We employ RKHS theory in the reformulation of conditional moment restriction (CMR) as a maximum moment restriction (MMR) based on which we can approach the problem from the empirical risk minimization (ERM) perspective. As we demonstrate, this framework not only facilitates theoretical analysis, but also results in easy-to-use algorithms that perform well in practice compared to existing methods. The paper also shows a way of combining the elegant theoretical approach of kernel methods with practical merits of deep neural networks. Despite these advantages, the optimal choice of the kernel $k$ remains an open question which we hope to address in future work.

In the next chapter, we present a conclusion summarizing the thesis and provide future directions.

## Conclusions and Future Work

In this chapter, we first summarize the contributions we made in the thesis. We then present some potential research directions for future work.

### 7.1 Conclusions

Given a set of data and the model hypothesis space, accurate estimation of the model parameters is crucial for e.g. data modeling and model-based inference. The thesis is motivated by the non-parametric Bayesian estimation of the Hawkes process. The nonparametric form provides powerful modeling ability and the approximate Bayesian inference approach is a natural choice for less sensitivity to randomness of finite samples and better scalability on large-scale datasets. For this end, the thesis studies two different kinds of efficient approximate inference frameworks, namely, the Laplace Bayesian Hawkes process (LBHP) as Chapter 3 and the Variational Bayesian Hawkes process (VBHP) as Chapter 4. Both frameworks employ the branching structure of the Hawkes process to simplify the Bayesian inference and the finite support assumption for acceleration. Notably, they rely on different approximate principles: LBHP exploits Gibbs sampling as the high-level procedure with Laplace approximation applied in each iteration; VBHP only exploits the variational inference method. As a result, VBHP has an advantage of efficiently selecting the hyper-parameters without grid search, while LBHP is a more general framework and is able to estimate the distribution of other quantities related to the Hawkes process, such as the branching factor.

Beyond the studies on the Hawkes process, the thesis further explores a new approximate Bayesian inference approach in Chapter 5, which is applicable to more general Gaussian process models. The approach builds on the expectation propagation algorithm and substitutes the $L_{2}$ Wasserstein distance to the Kullback-Leibler (KL) divergence. Interestingly, although equipped with a more complicated distance, the approach preserves some desirable properties, such as the locality property of the EP update.

Finally, the thesis studies a simple and robust frequentist estimation framework as Chapter 6, which is outside the Bayesian field. The framework is designed for estimation of conditional moment restriction models, including the Hawkes process, and depends on the kernel maximum moment restriction. The framework has a
simple risk function and eases the proofs of consistency and asymptotic normality properties. It can also result in efficient hyper-parameter selection for some specific model classes.

### 7.2 Future Research Directions

According to recent frequentist estimation methods of the Hawkes process as reviewed in Section 2.5, it would be interesting to develop efficient Bayesian versions for these methods to gain more robustness and better modeling performance. Some of these methods are based on normalizing flows, which can be applied for approximate Bayesian inference. Therefore, developing more practical normalizing flows is an interesting direction.

The kernel maximum moment restriction based method is general and we expect it to be applied in different subfields of stochastic point processes.

# Appendix: Laplace Bayesian Hawkes Process 

## A. 1 Computing the Integral Term of the Log-likelihood

We consider $\Omega=[0, T]$, the background intensity $\mu$, the triggering kernel $\phi(\cdot)=$ $1 / 2 f(\cdot)^{2}, f(\cdot)=\boldsymbol{\omega}^{T} \boldsymbol{e}(\cdot)$, and data $\left\{x_{i}\right\}_{i=1}^{N}$, and the integral term in the log-likelihood is calculated as below

$$
\begin{aligned}
\text { Integral Term } & =-\frac{1}{2} \sum_{i=1}^{N} \int_{0}^{T} f^{2}\left(t-x_{i}\right) d t \\
& =-\frac{1}{2} \sum_{i=1}^{N} \int_{0}^{T-x_{i}}\left[\sum_{k=1}^{K} \omega_{k} e_{k}(t)\right]^{2} d t \\
& =-\frac{1}{2} \sum_{i=1}^{N} \sum_{k=1}^{K} \sum_{k^{\prime}=1}^{K} \omega_{k} \omega_{k^{\prime}} \underbrace{\int_{0}^{T-x_{i}} e_{k}(t) e_{k^{\prime}}(t) d t}_{U_{k k^{\prime}}^{(i)}} \\
& =-\frac{1}{2} \sum_{i=1}^{N} \boldsymbol{\omega}^{T} U^{(i)} \boldsymbol{\omega} .
\end{aligned}
$$

In our case, Equation (3.9) has $d=1$, i.e., $\phi_{k}(x)=(2 / \pi)^{1 / 2} \sqrt{1 / 2}^{[k-1=0]} \cos [(k-$ 1) $x], k=1,2, \cdots$. The matrix $U^{(i)}$ is calculated as below:

$$
\begin{aligned}
U_{1,1}^{(i)} & =\int_{0}^{T-t_{1}} \frac{1}{\pi} d t=\frac{T-x_{i}}{\pi} \\
U_{k>1,1}^{(i)} & =U_{1, k>1}^{(i)}=\frac{\sqrt{2}}{\pi} \frac{\sin \left[(k-1)\left(t_{m}-x_{i}\right)\right]}{k-1} \\
U_{k, k(k>1)} & =\frac{1}{\pi}\left\{T-x_{i}+\frac{\sin \left[2(k-1)\left(T-x_{i}\right)\right]}{2(k-1)}\right\}, \\
U_{k, k^{\prime}\left(k \neq k^{\prime}\right)} & =\frac{1}{\pi}\left\{\frac{\sin \left[\left(k-k^{\prime}\right)\left(T-x_{i}\right)\right]}{k-k^{\prime}}+\frac{\sin \left[\left(k+k^{\prime}-2\right)\left(T-x_{i}\right)\right]}{k+k^{\prime}-2}\right\} .
\end{aligned}
$$

## A. 2 M.A.P. $\mu$ and $\phi$ Given Infinite Branching Structures

M.A.P. $\mu$ and $\phi$ given Infinite branching structures is written as:

$$
\begin{aligned}
& \underset{\boldsymbol{\omega}, \mu}{\operatorname{argmax}} \mathbb{E}_{B}\left[\log p\left(\boldsymbol{\omega}, \mu \mid B,\left\{x_{i}\right\}_{i=1}^{N}, \Omega, k\right)\right] \\
& =\underset{\boldsymbol{\omega}, \mu}{\operatorname{argmax}} \underbrace{\mathbb{E}_{B}\left[\log p\left(\left\{x_{i}\right\}_{i=1}^{N} \mid \boldsymbol{\omega}, \mu, B, \Omega, k\right)\right]}_{\text {Expected Log-likelihood }}+\underbrace{\log p(\boldsymbol{\omega})+\log p(\mu)}_{\text {Constraints }} \\
& =\underset{\boldsymbol{\omega}, \mu}{\operatorname{argmax}} \sum_{i=1}\left\{\sum_{x_{j}<x_{i}} p_{i j} \log \frac{1}{2}\left[\boldsymbol{\omega}^{T} \boldsymbol{e}\left(x_{i}-x_{j}\right)\right]^{2}-p_{i 0} \log \mu-\frac{1}{2} \int_{0}^{T-x_{i}}\left[\boldsymbol{\omega}^{T} \boldsymbol{e}(t)\right]^{2} d t\right\} \\
& \quad-(\beta+1) \mu T-\frac{1}{2} \boldsymbol{\omega}^{T} \Lambda^{-1} \boldsymbol{\omega}-(\alpha-1) \log \mu,
\end{aligned}
$$

where $B$ represents the branching structure, $p_{i j}$ the probabilities of triggering relationships shown as Equation (2.2), and $\alpha, \beta$ are parameters of the Gamma prior of $\mu$ T. The second line is obtained using Bayes' rule, which shows M.A.P. $\mu$ and $\phi$ given infinite branching structures is equivalent to maximizing the constrained expected log-likelihood, i.e., the objective function for the M-step of the EM algorithm and the third line is an explicit expression of the second line.

## A. 3 Mode-Finding the Triggering Kernel

Here we demonstrate in detail the computational challenges involved in finding the posterior mode with respect to the value of the triggering kernel at multiple point locations. Consider the triggering kernel $\phi(\cdot)=\frac{1}{2} f^{2}(\cdot)$ where $f(\cdot)$ is Gaussian process distributed. For a dataset $\left\{x_{i}\right\}_{i=1}^{N}, \boldsymbol{X} \equiv\left\{f\left(x_{i}\right)\right\}_{i=1}^{N}=\left\{X_{i}\right\}_{i=1}^{N}$ has a normal distribution, i.e., $\left\{f\left(x_{i}\right)\right\}_{i=1}^{N} \sim \mathcal{N}(\boldsymbol{m}, \boldsymbol{\Sigma})$ where $\boldsymbol{m}$ and $\boldsymbol{\Sigma}$ are the mean and the covariance matrix. The distribution of $Y \equiv\left\{\phi\left(x_{i}\right)\right\}_{i=1}^{N}=\left\{Y_{i}\right\}_{i=1}^{N}$ is derived as below where $F$ is the cumulative density function and $f$ the probabilistic density function.

$$
\begin{aligned}
& F_{Y}(\boldsymbol{y}) \\
& =P\left(-\sqrt{2 y_{i}}<X_{i}<\sqrt{2 y_{i}}, i=1, \cdots, N\right) \\
& =\int_{-\sqrt{2 y_{1}}}^{\sqrt{2 y_{1}}} \cdots \int_{-\sqrt{2 y_{N}}}^{\sqrt{2 y_{N}}} \frac{1}{\sqrt{(2 \pi)^{N} \boldsymbol{\Sigma}^{-1}}} \exp \left[-\frac{(\boldsymbol{X}-\boldsymbol{m})^{T} \boldsymbol{\Sigma}^{-1}(\boldsymbol{X}-\boldsymbol{m})}{2}\right] \mathrm{d} X_{1} \cdots \mathrm{~d} X_{N}, \\
& f_{Y}(\boldsymbol{y}) \\
& =\frac{\partial^{N}}{\partial y_{1} \cdots \partial y_{N}} F_{Y}(\boldsymbol{y}) \\
& =\frac{1}{\sqrt{(2 \pi)^{N} \boldsymbol{\Sigma}^{-1}}}\left(\Pi_{i=1}^{N} \frac{1}{2 \sqrt{2 y_{1}}}\right) \sum_{X \in X_{i=1}^{N}\left\{\sqrt{2 y_{1}},-\sqrt{2 y_{1}}\right\}} \exp \left[-\frac{(\boldsymbol{X}-\boldsymbol{m})^{T} \boldsymbol{\Sigma}^{-1}(\boldsymbol{X}-\boldsymbol{m})}{2}\right],
\end{aligned}
$$

where $\times$ is the Cartesian product. There are $2^{N}$ summations of exponential functions, which is intractable.


Figure A.1: Triggering kernels estimated by the Gibbs-Hawkes method (Section 3.2) and the EM-Hawkes method (Section 3.3.2). The true kernel is plotted as the bold gray curve. We plot the median (red) and [0.1, 0.9] interval (filled red) of the approximate predictive distribution, along with the triggering kernel inferred by the EM Hawkes method (blue). The hyper-parameters $a$ and $b$ of the Gaussian process kernel are set to 0.002 .

# Appendix: Variational Bayesian Hawkes Process 

## B. 1 Deriving Equation (4.5)

As per Equation (2.7), there is

$$
\begin{aligned}
& \operatorname{CELBO}(q(B, \boldsymbol{\mu}, f, \boldsymbol{u}), p(\mathcal{D} \mid B, \mu, f, \boldsymbol{u}), p(B, \mu, f, \boldsymbol{u})) \\
& =\mathbb{E}_{q(B, \mu, f)}[\log p(\mathcal{D} \mid B, \mu, f)]-\operatorname{KL}(q(B, \mu, f, \boldsymbol{u}) \| p(B, \mu, f, \boldsymbol{u}))
\end{aligned}
$$

where the KL term can be simplified as

$$
\begin{aligned}
& \mathrm{KL}(q(B, \mu, f, \boldsymbol{u}) \| p(B, \mu, f, \boldsymbol{u})) \\
& =\sum_{B} \iiint q(B, \mu, f, \boldsymbol{u}) \log \frac{q(B) q(\mu) p(f \mid \boldsymbol{u}) q(\boldsymbol{u})}{p(B) p(\mu) p(f \mid \boldsymbol{u}) p(\boldsymbol{u})} \mathrm{d} \boldsymbol{u} \mathrm{~d} f \mathrm{~d} \mu \quad \text { (Equation (4.4) and Bayes' rule) } \\
& =\sum_{B} \iiint q(B, \mu, f, \boldsymbol{u}) \mathrm{d} \boldsymbol{u} \mathrm{~d} f \mathrm{~d} \mu \log \frac{q(B)}{p(B)} \\
& \quad+\sum_{B} \iiint q(B, \mu, f, \boldsymbol{u}) \mathrm{d} \boldsymbol{u} \mathrm{~d} f \log \frac{q(\mu)}{p(\mu)} \mathrm{d} \mu \\
& \quad+\sum_{B} \iiint q(B, \mu, f, \boldsymbol{u}) \mathrm{d} f \mathrm{~d} \mu \log \frac{q(\boldsymbol{u})}{p(\boldsymbol{u})} \mathrm{d} \boldsymbol{u} \quad \text { (simplification) } \\
& = \\
& =\sum_{B} q(B) \log \frac{q(B)}{p(B)}+\int q(\mu) \log \frac{q(\mu)}{p(\mu)} \mathrm{d} \mu+\int q(\boldsymbol{u}) \log \frac{q(\boldsymbol{u})}{p(\boldsymbol{u})} \mathrm{d} \boldsymbol{u} \quad \text { (simplification) } \\
& =
\end{aligned}
$$

We utilise the likelihood $p(\mathcal{D}, B \mid \mu, f)$ by the reconstruction term and the KL term w.r.t. B

$$
\begin{aligned}
& \mathbb{E}_{q(B, \mu, f)}[\log p(\mathcal{D} \mid B, \mu, f)]-\operatorname{KL}(q(B) \| p(B)) \\
& =\sum_{B} \iint q(B, \mu, f) \log p(\mathcal{D} \mid B, \mu, f) \mathrm{d} f \mathrm{~d} \mu-\sum_{B} q(B) \log \frac{q(B)}{p(B)} \quad \text { (definition) }
\end{aligned}
$$

$$
\begin{aligned}
& =\sum_{B} \iint q(B, \mu, f) \log p(\mathcal{D} \mid B, \mu, f) \mathrm{d} f \mathrm{~d} \mu-\sum_{B} \iint q(B, \mu, f) \log \frac{q(B)}{p(B)} \mathrm{d} f \mathrm{~d} \mu \\
& \quad \text { (align probabilities) } \\
& =\sum_{B} \iint q(B, \mu, f) \log \frac{p(\mathcal{D} \mid B, \mu, f) p(B)}{q(B)} \mathrm{d} f \mathrm{~d} \mu \quad \text { (merge) } \\
& =\sum_{B} \iint q(B, \mu, f) \log p(\mathcal{D}, B \mid \mu, f) \mathrm{d} f \mathrm{~d} \mu-\sum_{B} q(B) \log q(B) \quad \text { (merge) } \\
& =\mathbb{E}_{q(B, \mu, f)}[\log p(\mathcal{D}, B \mid \mu, f)]+H_{B}
\end{aligned}
$$

where $H_{B}=-\sum_{B} q(B) \log q(B)$ is the entropy of $B$ and further computed as follows. We adopt $q(B)$ from Equation (4.3) and the close form expression of $H_{B}$ is derived as

$$
\begin{aligned}
H_{B} & =-\sum_{\left\{\boldsymbol{b}_{i}\right\}_{i=1}^{N}}\left(\prod_{i=1}^{N} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right) \log \left(\prod_{i=1}^{N} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right) \\
& =-\sum_{j=0}^{k-1} \sum_{\left\{\boldsymbol{b}_{i}\right\}_{i \neq k}}\left(q_{k j} \prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right) \log \left(q_{k j} \prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right) \quad \text { (split the summation) } \\
& =-\sum_{j=0}^{k-1} \sum_{\left\{\boldsymbol{b}_{i}\right\}_{i \neq k}}\left(q_{k j} \prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right)\left[\log q_{k j}+\log \left(\prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right)\right] \quad \text { (split the log) } \\
& =-\sum_{j=0}^{k-1} q_{k j} \log q_{k j} \underbrace{\sum_{\left\{\boldsymbol{b}_{i}\right\}_{i \neq k}}\left(\prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right)}_{=1}-\underbrace{\sum_{j=0}^{k-1} q_{k j} \sum_{\left\{\boldsymbol{b}_{i}\right\}_{i \neq k}}\left(\prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right) \log \left(\prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right)}_{=1}
\end{aligned}
$$

(distributive law of multiplication)
$=-\sum_{j=0}^{k-1} q_{k j} \log q_{k j}-\sum_{\left\{\boldsymbol{b}_{i}\right\}_{i \neq k}}\left(\prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right) \log \left(\prod_{i \neq k} \prod_{j=0}^{i-1} q_{i j}^{b_{i j}}\right) \quad$ (simplification)
$=\cdots \quad$ (do the same operations as above)
$=-\sum_{i=1}^{N} \sum_{j=0}^{i-1} q_{i j} \log q_{i j}$.

## B. 2 Closed Forms of KL Terms in the ELBO

$$
\begin{aligned}
& \operatorname{KL}(q(\mu) \| p(\mu))=\left(k-k_{0}\right) \psi(k)-k_{0} \log \left(c / c_{0}\right)-k-\log \left[\Gamma(k) / \Gamma\left(k_{0}\right)\right]+c k / c_{0} \\
& \operatorname{KL}(q(\boldsymbol{u}) \| p(\boldsymbol{u}))=\left[\operatorname{Tr}\left(\boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{S}\right)+\log \left|\boldsymbol{K}_{z z^{\prime}}\right| /|\boldsymbol{S}|-M+\boldsymbol{m}^{T} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{m}\right] / 2
\end{aligned}
$$

## B. 3 Extra Closed Form Expressions for Equation (4.2.2)

$$
\begin{aligned}
& \int_{\mathcal{T}_{i}} \mathbb{E}_{q(f)}^{2}(f)=\int_{\mathcal{T}_{i}} \boldsymbol{m}^{T} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{K}_{z x} \boldsymbol{K}_{x z} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{m} \mathrm{~d} \boldsymbol{x} \\
& =\boldsymbol{m}^{T} \boldsymbol{K}_{z z^{\prime}}^{-1}\left(\int_{\mathcal{T}_{i}} \boldsymbol{K}_{z x} \boldsymbol{K}_{x z} \mathrm{~d} x\right) \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{m} \\
& \equiv \boldsymbol{m}^{T} \boldsymbol{K}_{z z^{\prime}}^{-1} \mathbf{\Psi}_{i} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{m} \text {. } \\
& \int_{\mathcal{T}_{i}} \operatorname{Var}_{q(f)}[f]=\int_{\mathcal{T}_{i}} \boldsymbol{K}_{x x}-\boldsymbol{K}_{x z} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{K}_{z x}+\boldsymbol{K}_{x z} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{S} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{K}_{z x} \mathrm{~d} \boldsymbol{x} \\
& =\int_{\mathcal{T}_{i}} \gamma-\operatorname{Tr}\left(\boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{K}_{z x} \boldsymbol{K}_{x z}\right)+\operatorname{Tr}\left(\boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{S} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{K}_{z x} \boldsymbol{K}_{x z}\right) \mathrm{d} \boldsymbol{x} \\
& =\int_{\mathcal{T}_{i}} \gamma \mathrm{~d} \boldsymbol{x}-\operatorname{Tr}\left(\boldsymbol{K}_{z z^{\prime}}^{-1} \int_{\mathcal{T}_{i}} \boldsymbol{K}_{z x} \boldsymbol{K}_{x z} \mathrm{~d} \boldsymbol{x}\right)+\operatorname{Tr}\left(\boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{S} \boldsymbol{K}_{z z^{\prime}}^{-1} \int_{\mathcal{T}_{i}} \boldsymbol{K}_{z x} \boldsymbol{K}_{x z} \mathrm{~d} \boldsymbol{x}\right) \\
& \equiv \gamma\left|\mathcal{T}_{i}\right|-\operatorname{Tr}\left(\boldsymbol{K}_{z z^{\prime}}^{-1} \mathbf{\Psi}_{i}\right)+\operatorname{Tr}\left(\boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{S} \boldsymbol{K}_{z z^{\prime}}^{-1} \boldsymbol{\Psi}_{i}\right) . \\
& \boldsymbol{\Psi}_{i}\left(\boldsymbol{z}, \boldsymbol{z}^{\prime}\right)=\int_{\mathcal{T}_{i}} \boldsymbol{K}_{z x} \boldsymbol{K}_{x z^{\prime}} \mathrm{d} x \\
& =\int_{\mathcal{T}_{i}} \gamma^{2} \prod_{r=1}^{R} \exp \left(-\frac{\left(x_{r}-z_{r}\right)^{2}}{2 \alpha_{r}}\right) \exp \left(-\frac{\left(z_{r}^{\prime}-x_{r}\right)^{2}}{2 \alpha_{r}}\right) \mathrm{d} x \\
& =\int_{\mathcal{T}_{i}} \gamma^{2} \prod_{r=1}^{R} \exp \left(-\frac{\left(z_{r}-z_{r}^{\prime}\right)^{2}}{4 \alpha_{r}}\right) \exp \left(-\frac{\left(\bar{z}_{r}-x_{r}\right)^{2}}{\alpha_{r}}\right) \mathrm{d} x \\
& =\gamma^{2} \prod_{r=1}^{R} \exp \left(-\frac{\left(z_{r}-z_{r}^{\prime}\right)^{2}}{4 \alpha_{r}}\right) \int_{\mathcal{T}_{i, r}} \exp \left(-\frac{\left(\bar{z}_{r}-x_{r}\right)^{2}}{\alpha_{r}}\right) \mathrm{d} x_{r} \quad\left(y_{r}=\left(\bar{z}_{r}-x_{r}\right) / \alpha_{r}\right) \\
& =\gamma^{2} \prod_{r=1}^{R}-\sqrt{\alpha_{r}} \exp \left(-\frac{\left(z_{r}-z_{r}^{\prime}\right)^{2}}{4 \alpha_{r}}\right) \int_{\left(\bar{z}_{r}-\mathcal{T}_{i, r}^{\text {min }}\right) / \sqrt{\alpha_{r}}}^{\left(\bar{z}_{r}-\mathcal{T}_{r}^{\text {max }}\right) / \sqrt{\alpha_{r}}} \exp \left(-y_{r}^{2}\right) \mathrm{d} y_{r} \\
& =\gamma^{2} \prod_{r=1}^{R}-\frac{\sqrt{\pi \alpha_{r}}}{2} \exp \left(-\frac{\left(z_{r}-z_{r}^{\prime}\right)^{2}}{4 \alpha_{r}}\right)\left[\operatorname{erf}\left(\frac{\bar{z}_{r}-\mathcal{T}_{i, r}^{\text {max }}}{\sqrt{\alpha_{r}}}\right)-\operatorname{erf}\left(\frac{\bar{z}_{r}-\mathcal{T}_{i, r}^{\text {min }}}{\sqrt{\alpha_{r}}}\right)\right] \text {. }
\end{aligned}
$$

where we explicitly express $\mathcal{T}_{i}$ as a Cartesian product $\mathcal{T}_{i} \equiv X_{r=1}^{R}\left[\mathcal{T}_{i, r}^{\min }, \mathcal{T}_{i, r}^{\max }\right]$ for $R$-dimensional data.

## B. 4 Additional Experiment Results



Figure B.1: Posterior Triggering Kernels Inferred By VBHP and Gibbs Hawkes. Results of Gibbs Hawkes are obtained in 2000 iterations.


Figure B.2: Convergence Rate of VBHP and Gibbs Hawkes with Different Numbers of Inducing Points. VBHP and Gibbs Hawkes measure respectively the relative error of the approximate marginal likelihood and of the posterior distribution of the Gaussian process.

## Appendix: Quantile Propagation

## C. 1 Minimization of $L_{2}$ WD between Univariate Gaussian and Non-Gaussian Distributions

In this section, we derive the formulas of the optimal $\mu^{*}$ and $\sigma^{*}$ for the $L_{2} \mathrm{WD}$, i.e., , Equation (5.1). Recall the optimization problem: we use a univariate Gaussian distribution $\mathcal{N}\left(f \mid \mu, \sigma^{2}\right)$ to approximate a univariate non-Gaussian distribution $q(f)$ by minimizing the $L_{2}$ WD between them:

$$
\min _{\mu, \sigma} W_{2}^{2}(q, \mathcal{N})=\min _{\mu, \sigma} \int_{0}^{1}\left|F_{q}^{-1}(y)-\mu-\sqrt{2} \sigma \operatorname{erf}^{-1}(2 y-1)\right|^{2} \mathrm{~d} y,
$$

where $F_{q}^{-1}$ is the quantile function of the non-Gaussian distribution $q$, namely the pseudo-inverse function of the corresponding cumulative distribution function $F_{q}$ defined in Proposition 1.

To solve this problem, we first calculate derivatives about $\mu$ and $\sigma$ :

$$
\begin{aligned}
\frac{\partial \mathrm{W}_{2}^{2}}{\partial \mu} & =-2 \int_{0}^{1} F_{q}^{-1}(y)-\mu-\sqrt{2} \sigma \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y, \\
\frac{\partial \mathrm{~W}_{2}^{2}}{\partial \sigma} & =-2 \int_{0}^{1}\left(F_{q}^{-1}(y)-\mu-\sqrt{2} \sigma \operatorname{erf}^{-1}(2 y-1)\right) \sqrt{2} \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y .
\end{aligned}
$$

Then, by zeroing derivatives, we obtain the optimal parameters:

$$
\begin{aligned}
\mu^{*} & =\int_{0}^{1} F_{q}^{-1}(y)-\sqrt{2} \sigma \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y \\
& =\int_{-\infty}^{\infty} x q(x) \mathrm{d} x-\frac{\sqrt{2}}{2} \sigma \int_{-1}^{1} \operatorname{erf}^{-1}(y) \mathrm{d} y \\
& =\mu_{q}-\sqrt{2} \sigma \int_{-\infty}^{\infty} x \mathcal{N}(x \mid 0,1 / 2) \mathrm{d} x \\
& =\mu_{q} \\
\sigma^{*} & =\sqrt{2} \int_{0}^{1}\left(F_{q}^{-1}(y)-\mu\right) \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y / \int_{0}^{1} 2\left(\operatorname{erf}^{-1}\right)^{2}(2 y-1) \mathrm{d} y
\end{aligned}
$$

$$
\begin{align*}
& =\sqrt{2} \int_{0}^{1} F_{q}^{-1}(y) \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y / \underbrace{\int_{-\infty}^{\infty} 2 x^{2} \mathcal{N}(x \mid 0,1 / 2) \mathrm{d} x}_{=1} \\
& =\sqrt{2} \int_{0}^{1} F_{q}^{-1}(y) \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y \\
& =\sqrt{2} \int_{-\infty}^{\infty} f \operatorname{erf}^{-1}\left(2 F_{q}(f)-1\right) \mathrm{d} F_{q}(f) \\
& =-\sqrt{\frac{1}{2 \pi}} \int_{-\infty}^{\infty} f \mathrm{~d} e^{-\left[\operatorname{erf}^{-1}\left(2 F_{\bar{q}}(f)-1\right)\right]^{2}} \\
& =0+\sqrt{\frac{1}{2 \pi}} \int_{-\infty}^{\infty} e^{-\left[\operatorname{erf}^{-1}\left(2 F_{\bar{q}}(f)-1\right)\right]^{2}} \mathrm{~d} f . \tag{C.1}
\end{align*}
$$

## C. 2 Minimization of $L_{p}$ WD between Univariate Gaussian and Non-Gaussian Distributions

In this section, we describe a gradient descent approach to minimizing an $L_{p} \mathrm{WD}$, for $p \neq 2$, in order to handle cases with no analytical expressions for the optimal parameters. Our goal is to use a univariate Gaussian distribution $\mathcal{N}\left(f \mid \mu, \sigma^{2}\right)$ to approximate a univariate non-Gaussian distribution $q(f)$. Specifically, we seek the minimizer in $\mu$ and $\sigma$ of $W_{p}^{p}(q, \mathcal{N})$; the derivatives of the objective function about $\mu$ and $\sigma$ are:

$$
\begin{aligned}
\partial_{\mu} \mathrm{W}_{p}^{p} & =-p \int_{0}^{1}|\varepsilon(y)|^{p-1} \operatorname{sgn}(\varepsilon(y)) \mathrm{d} y=-p \int_{-\infty}^{\infty}|\eta(x)|^{p-1} \operatorname{sgn}(\eta(x)) q(x) \mathrm{d} x, \\
\partial_{\sigma} \mathrm{W}_{p}^{p} & =-p \int_{0}^{1}|\varepsilon(y)|^{p-1} \operatorname{sgn}(\varepsilon(y)) \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y \\
& =-p \int_{-\infty}^{\infty}|\eta(x)|^{p-1} \operatorname{sgn}(\eta(x)) \operatorname{erf}^{-1}\left(2 F_{q}(x)-1\right) q(x) \mathrm{d} x .
\end{aligned}
$$

where for simplification, we define $\varepsilon(y)=F_{q}^{-1}(y)-\mu-\sqrt{2} \sigma \operatorname{erf}^{-1}(2 y-1)$ and $\eta(x)=$ $x-\mu-\sqrt{2} \sigma \operatorname{erf}^{-1}\left(2 F_{q}(x)-1\right)$, with $F_{q}$ and $F_{q}^{-1}$ being the CDF and the quantile function of $q$. Note the derivatives have no analytical expressions. However, if the $\operatorname{CDF} F_{q}$ is available, we can use the standard numerical integration routines; otherwise, we resort to Monte Carlo sampling. In the framework of EP or QP, $q(x) \propto q^{\backslash i}(x) p\left(y_{i} \mid x\right)$ and $q^{\backslash i}$ is Gaussian, so we may draw samples from a Gaussian proposal distribution to obtain a simple Monte Carlo method.

## C. 3 Computations for Different Likelihoods

Given the likelihood $p(y \mid f)$ and the cavity distribution $q^{i}(f)=\mathcal{N}\left(f \mid \mu, \sigma^{2}\right)$, a stable way to compute the mean and the variance of the tilted distribution $\widetilde{q}(f)=$ $p(y \mid f) q^{i}(f) / Z$ where the normalizer $Z=\int_{-\infty}^{\infty} p(y \mid f) q^{i}(f) \mathrm{d} f$, can be found in the software manual of Rasmussen and Williams [2005]. We present the key formulae
below, for use in subsequent derivations:

$$
\begin{aligned}
\partial_{\mu} Z & =\int_{-\infty}^{\infty} \frac{f-\mu}{\sigma^{2}} p(y \mid f) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right) \mathrm{d} f \\
\frac{\partial_{\mu} Z}{Z} & =\frac{1}{\sigma^{2}} \int_{-\infty}^{\infty} f \frac{p(y \mid f) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right)}{Z} \mathrm{~d} f-\frac{\mu}{\sigma^{2}} \int_{-\infty}^{\infty} \frac{p(y \mid f) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right)}{Z} \mathrm{~d} y \\
\frac{\partial_{\mu} Z}{Z} & =\frac{1}{\sigma^{2}} \mu_{\tilde{q}}-\frac{\mu}{\sigma^{2}} \\
\Longrightarrow \mu_{\tilde{q}} & =\frac{\sigma^{2} \partial_{\mu} Z}{Z}+\mu=\sigma^{2} \partial_{\mu} \log Z+\mu, \\
\partial_{\mu}^{2} Z & =\int_{-\infty}^{\infty}-\frac{1}{\sigma^{2}} p(y \mid f) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right)+\left(\frac{f-\mu}{\sigma^{2}}\right)^{2} p(y \mid f) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right) \mathrm{d} f \\
\frac{\partial_{\mu}^{2} Z}{Z} & =\int_{-\infty}^{\infty}\left(-\frac{1}{\sigma^{2}}+\frac{\mu^{2}}{\sigma^{4}}+\frac{f^{2}}{\sigma^{4}}-\frac{2 \mu f}{\sigma^{4}}\right) \frac{p(y \mid f) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right)}{Z} \mathrm{~d} f \\
\frac{\partial_{\mu}^{2} Z}{Z} & =-\frac{1}{\sigma^{2}}+\frac{\mu^{2}}{\sigma^{4}}+\frac{1}{\sigma^{4}}\left(\sigma_{\tilde{q}}^{2}+\mu_{\tilde{q}}^{2}\right)-\frac{2 \mu}{\sigma^{4}} \mu_{\tilde{q}} \\
\frac{\partial_{\mu}^{2} Z}{Z} & =-\frac{1}{\sigma^{2}}+\frac{\sigma_{\tilde{q}}^{2}}{\sigma^{4}}+\frac{\left(\mu-\mu_{\tilde{q}}\right)^{2}}{\sigma^{4}}=-\frac{1}{\sigma^{2}}+\frac{\sigma_{\tilde{q}}^{2}}{\sigma^{4}}+\left(\frac{\partial_{\mu} Z}{Z}\right)^{2} \\
\Longrightarrow \sigma_{\tilde{q}}^{2} & =\sigma^{4}\left[\frac{\partial_{\mu}^{2} Z}{Z}-\left(\frac{\partial_{\mu} Z}{Z}\right)^{2}\right]+\sigma^{2}=\sigma^{4} \partial_{\mu}^{2} \log Z+\sigma^{2} .
\end{aligned}
$$

## C.3.1 Probit Likelihood for Binary Classification

For the binary classification with labels $y \in\{-1,1\}$, the PDF of the tilted distribution $\widetilde{q}(f)$ with the probit likelihood is provided by Rasmussen and Williams [2005]:

$$
\widetilde{q}(f)=Z^{-1} \Phi(f y) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right), \quad Z=\Phi(z), \quad z=\frac{\mu}{y \sqrt{1+\sigma^{2}}}
$$

and the mean estimate also has a closed form expression:

$$
\mu^{\star}=\mu_{\widetilde{q}}=\mu+\frac{\sigma^{2} \mathcal{N}(z)}{\Phi(z) y \sqrt{1+\sigma^{2}}} .
$$

As per Equation (5.1), the computation of the optimal $\sigma^{\star}$ requires the CDF of $\widetilde{q}$, denoted as $F_{\tilde{q}}$. For positive $y>0$, the CDF is derived as

$$
\begin{aligned}
F_{\widetilde{q}, y>0}(x) & =Z^{-1} \int_{-\infty}^{x} \Phi(f y) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right) \mathrm{d} f \\
& =\frac{Z^{-1}}{2 \pi \sigma y} \int_{-\infty}^{\mu} \int_{-\infty}^{x-\mu} \exp \left(-\frac{1}{2}\left[\begin{array}{c}
w \\
f
\end{array}\right]^{T}\left[\begin{array}{cc}
v^{-2}+\sigma^{-2} & v^{-2} \\
v^{-2} & v^{-2}
\end{array}\right]\left[\begin{array}{c}
w \\
f
\end{array}\right]\right) \mathrm{d} w \mathrm{~d} f \\
& =Z^{-1} \int_{-\infty}^{k} \int_{-\infty}^{h} \mathcal{N}\left(\left.\left[\begin{array}{c}
w \\
f
\end{array}\right] \right\rvert\, \mathbf{0},\left[\begin{array}{cc}
1 & -\rho \\
-\rho & 1
\end{array}\right]\right) \mathrm{d} w \mathrm{~d} f
\end{aligned}
$$

$$
\begin{aligned}
& \stackrel{(a)}{=} Z^{-1}\left[\frac{1}{2} \Phi(h)-T\left(h, \frac{k+\rho h}{h \sqrt{1-\rho^{2}}}\right)+\frac{1}{2} \Phi(k)-T\left(k, \frac{h+\rho k}{k \sqrt{1-\rho^{2}}}\right)+\eta\right] \\
& k=\frac{\mu}{\sqrt{\sigma^{2}+1}}, \quad h=\frac{x-\mu}{\sigma}, \quad \rho=\frac{1}{\sqrt{1+1 / \sigma^{2}}}, \quad x \neq \mu, \quad \mu \neq 0,
\end{aligned}
$$

where the step (a) is obtained by exploiting the work of Owen [1956] and $T(\cdot, \cdot)$ is the Owen's T function:

$$
T(h, a)=\frac{1}{2 \pi} \int_{0}^{a} \frac{\exp \left[-\left(1+x^{2}\right) h^{2} / 2\right]}{1+x^{2}} \mathrm{~d} x,
$$

and $\eta$ is defined as

$$
\eta= \begin{cases}0 & h k>0 \text { or }(h k=0 \text { and } h+k \geq 0) \\ -0.5 & \text { otherwise }\end{cases}
$$

Similarly, for $y<0$, the CDF is

$$
F_{\tilde{q}, y<0}(x)=Z^{-1}\left[\frac{1}{2} \Phi(h)+T\left(h, \frac{k+\rho h}{h \sqrt{1-\rho^{2}}}\right)-\frac{1}{2} \Phi(k)+T\left(k, \frac{h+\rho k}{k \sqrt{1-\rho^{2}}}\right)-\eta\right] .
$$

Summarizing the two cases, we get the closed form expression of $F_{\overparen{q}}$ :

$$
\begin{aligned}
& F_{\tilde{\eta}}(x) \\
& =Z^{-1}\left[\frac{1}{2} \Phi(h)-y T\left(h, \frac{k+\rho h}{h \sqrt{1-\rho^{2}}}\right)+\frac{y}{2} \Phi(k)-y T\left(k, \frac{h+\rho k}{k \sqrt{1-\rho^{2}}}\right)+y \eta\right] \\
& =Z^{-1}\left[\frac{1}{2} \Phi(h)-y T\left(h, \frac{k}{h \sqrt{1-\rho^{2}}}+\sigma\right)+\frac{y}{2} \Phi(k)-y T\left(k, \frac{h}{k \sqrt{1-\rho^{2}}}+\sigma\right)+y \eta\right] .
\end{aligned}
$$

Provided the above, the optimal $\sigma^{\star}$ can be computed by numerical integration of Eqn (C.1). For special cases, we provide additional formulas:

$$
\begin{aligned}
& \text { (1) } x=\mu, \mu \neq 0: F_{\widetilde{q}}(x)=Z^{-1}\left[\frac{1}{4}-\frac{y \operatorname{sign}(k)}{4}+\frac{y}{2} \Phi(k)-y T(k, \sigma)+y \eta\right] ; \\
& \text { (2) } x \neq \mu, \mu=0: F_{\widetilde{\eta}}(x)=2\left[\frac{1}{2} \Phi(h)-y T(h, \sigma)+\frac{y}{4}-\frac{y \operatorname{sign}(h)}{4}+y \eta\right] ; \\
& \text { (3) } x=\mu, \mu=0: F_{\widetilde{q}}(x)=\frac{1}{2}-\frac{y}{\pi} \arctan (\sigma) .
\end{aligned}
$$

## C.3.2 Square Link Function for Poisson Regression

Consider Poisson regression, which uses the Poisson likelihood $p(y \mid g)=g^{y} \exp (-g) / y$ ! to model count data $y \in \mathbb{N}$, with the square link function $g(f)=f^{2}$ [Walder and

Bishop, 2017; Flaxman et al., 2017]. We use the square link function because it is more mathematically convenient than the exponential function. Given the cavity distribution $q^{\backslash i}(f)=\mathcal{N}\left(f \mid \mu, \sigma^{2}\right)$, we want the tilted distribution $\widetilde{q}(f)=q^{\backslash i}(f) p(y \mid g(f)) / Z$ where the normalizer Z is derived as:

$$
\begin{align*}
Z & =\int_{-\infty}^{\infty} q^{i i}(f) p(y \mid g) \mathrm{d} f \\
& =\int_{-\infty}^{\infty} \frac{1}{\sqrt{2 \pi \sigma^{2}}} \exp \left(-\frac{(f-\mu)^{2}}{2 \sigma^{2}}\right) f^{2 y} \exp \left(-f^{2}\right) / y!\mathrm{d} f \\
& \stackrel{(a)}{=} \frac{1}{\sqrt{2 \pi \sigma^{2}} y!\exp \left(\mu^{2} /\left(1+2 \sigma^{2}\right)\right)} \int_{-\infty}^{\infty} f^{2 y} \exp \left(-\frac{\left(f-\mu /\left(1+2 \sigma^{2}\right)\right)^{2}}{2 \sigma^{2} /\left(1+2 \sigma^{2}\right)}\right) \mathrm{d} f \\
& \stackrel{(\text { b) }}{=} \frac{\left(\frac{2 \sigma^{2}}{1+2 \sigma^{2}}\right)^{y+\frac{1}{2}}}{\sqrt{2 \pi \sigma^{2}} y!\exp \left(\mu^{2} /\left(1+2 \sigma^{2}\right)\right)} \Gamma\left(y+\frac{1}{2}\right){ }_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{\mu^{2}}{2 \sigma^{2}\left(1+2 \sigma^{2}\right)}\right) \\
& =\frac{\alpha^{y+\frac{1}{2}}}{\sqrt{2 \pi \sigma^{2}} y!\exp (h)} \Gamma\left(y+\frac{1}{2}\right){ }_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{h}{2 \sigma^{2}}\right), \\
\alpha & =\frac{2 \sigma^{2}}{1+2 \sigma^{2}}, \quad h=\frac{\mu^{2}}{1+2 \sigma^{2}} \tag{C.2}
\end{align*}
$$

where the step (a) rewrites the product of two exponential functions into the form of the Gaussian distribution, (b) is achieved through Mathematica [Wolfram, 2019], $\Gamma(\cdot)$ is the Gamma function and ${ }_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{h^{2}}{2 \sigma^{2}}\right)$ is the confluent hypergeometric function of the first kind. Furthermore, we compute the first derivative of $\log Z$ w.r.t. $\mu$ and then the mean of the tilted distribution:

$$
\begin{aligned}
\partial_{\mu} \log Z= & \left(\frac{y_{1} F_{1}\left(-y+1 ; \frac{3}{2} ;-\frac{h}{2 \sigma^{2}}\right)}{\sigma^{2}{ }_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{h}{2 \sigma^{2}}\right)}-1\right) \frac{2 \mu}{1+2 \sigma^{2}} \\
\Longrightarrow \mu_{\tilde{q}}= & \sigma^{2} \partial_{\mu} \log Z+\mu . \\
\partial_{\mu}^{2} \log Z= & \left(\frac{y_{1} F_{1}\left(-y+1 ; \frac{3}{2} ;-\frac{h}{2 \sigma^{2}}\right)}{\sigma^{2}{ }_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{h}{2 \sigma^{2}}\right)}-1\right) \frac{2}{1+2 \sigma^{2}}- \\
& \left(\frac{2(1-y)_{1} F_{1}\left(-y+2 ; \frac{5}{2} ;-\frac{h}{2 \sigma^{2}}\right)}{3_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{h}{2 \sigma^{2}}\right)}+\frac{2 y_{1} F_{1}\left(-y+1 ; \frac{3}{2} ;-\frac{h}{2 \sigma^{2}}\right)^{2}}{{ }_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{h}{2 \sigma^{2}}\right)^{2}}\right) \frac{2 \mu^{2} y}{\sigma^{4}\left(1+2 \sigma^{2}\right)^{2}} \\
\Longrightarrow \sigma_{\tilde{q}}^{2}= & \sigma^{4} \partial_{\mu}^{2} \log Z+\sigma^{2}
\end{aligned}
$$

Finally, we derive the CDF of the tilted distribution $\widetilde{q}$ by using the binomial theorem:

$$
F_{\widetilde{q}}(x)=Z^{-1} \int_{-\infty}^{x} p(y \mid g) \mathcal{N}\left(f \mid \mu, \sigma^{2}\right) \mathrm{d} f
$$

$$
\begin{aligned}
& \stackrel{(\text { (a) }}{=} A \int_{-\infty}^{x} f^{2 y} \exp \left(-\frac{\left(f-\mu /\left(1+2 \sigma^{2}\right)\right)^{2}}{2 \sigma^{2} /\left(1+2 \sigma^{2}\right)}\right) \mathrm{d} f \\
& =A \int_{-\infty}^{x-\frac{\mu}{1+2 \sigma^{2}}}\left(f+\frac{\mu}{1+2 \sigma^{2}}\right)^{2 y} \exp \left(-\frac{f^{2}}{2 \sigma^{2} /\left(1+2 \sigma^{2}\right)}\right) \mathrm{d} f \\
& \stackrel{\text { (b) }}{=} A \int_{-\infty}^{x-\beta}\left[\sum_{k=0}^{2 y}\binom{2 y}{k} f^{k} \beta^{2 y-k}\right] \exp \left(-\frac{f^{2}}{\alpha}\right) \mathrm{d} f \\
& =A \sum_{k=0}^{2 y}\binom{2 y}{k} \beta^{2 y-k}\left[\int_{-\infty}^{0} f^{k} \exp \left(-\frac{f^{2}}{\alpha}\right) \mathrm{d} f+\int_{0}^{x-\beta} f^{k} \exp \left(-\frac{f^{2}}{\alpha}\right) \mathrm{d} f\right] \\
& \stackrel{\text { (c) }}{=} \frac{A}{2} \sum_{k=0}^{2 y}\binom{2 y}{k} \beta^{2 y-k} \alpha^{\frac{k+1}{2}}\left[(-1)^{k} \Gamma\left(\frac{k+1}{2}\right)+\right. \\
& A=\frac{\left.\operatorname{sgn}(x-\beta)^{k+1}\left(\Gamma\left(\frac{k+1}{2}\right)-\Gamma\left(\frac{k+1}{2}, \frac{(x-\beta)^{2}}{\alpha}\right)\right)\right]}{\sqrt{2 \pi \sigma^{2}} y!\exp \left(\mu^{2} /\left(1+2 \sigma^{2}\right)\right)}=\left[\alpha^{y+\frac{1}{2}} \Gamma\left(y+\frac{1}{2}\right){ }_{1} F_{1}\left(-y ; \frac{1}{2} ;-\frac{h}{2 \sigma^{2}}\right)\right]^{-1}, \\
& \beta=\frac{\mu}{1+2 \sigma^{2}},
\end{aligned}
$$

where the step (a) has been derived in (a) of Equation (C.2), (b) applies the binomial theorem and (c) is obtained through Mathematica [Wolfram, 2019]. And, the function $\Gamma(a, z)=\int_{z}^{\infty} t^{a-1} e^{-t} \mathrm{~d} t$ is the upper incomplete gamma function and $\operatorname{sgn}(x)$ is the sign function, equaling 1 when $x>0,0$ when $x=0$ and -1 when $x<0$.

## C. 4 Proof of Convexity

Theorem. Given two probability measures in $\mathcal{M}_{+}^{1}(\mathbb{R})$ : a Gaussian $\mathcal{N}\left(\mu, \sigma^{2}\right)$ with mean $\mu$ and standard deviation $\sigma>0$, and an arbitrary measure $\widetilde{q}$, the $L_{p}$ WD $\mathrm{W}_{p}^{p}(\widetilde{q}, \mathcal{N})$ is strictly convex about $\mu$ and $\sigma$.

Proof. Let $F_{\tilde{q}}^{-1}(y)$ and $F_{\mathcal{N}}^{-1}(y)=\mu+\sqrt{2} \sigma \operatorname{erf}^{-1}(2 y-1), y \in[0,1]$, be the quantile functions of $\widetilde{q}$ and the Gaussian $\mathcal{N}$, where erf is the error function. Then, we consider two distinct Gaussian measures $\mathcal{N}\left(\mu_{1}, \sigma_{1}^{2}\right)$ and $\mathcal{N}\left(\mu_{2}, \sigma_{2}^{2}\right)$ and a convex combination w.r.t. their parameters $\mathcal{N}\left(a_{1} \mu_{1}+a_{2} \mu_{2},\left(a_{1} \sigma_{1}+a_{2} \sigma_{2}\right)^{2}\right)$ with $a_{1}, a_{2} \in \mathbb{R}_{+}$and $a_{1}+a_{2}=1$. Given the above, we further define $\varepsilon_{k}(y)=F_{\tilde{q}}^{-1}(y)-\mu_{k}-\sigma_{k} \sqrt{2} \operatorname{erf}^{-1}(2 y-1), k=1,2$, for notational simplification, and derive the convexity as:

$$
\begin{gathered}
\mathrm{W}_{p}^{p}\left(\widetilde{q}, \mathcal{N}\left(a_{1} \mu_{1}+a_{2} \mu_{2},\left(a_{1} \sigma_{1}+a_{2} \sigma_{2}\right)^{2}\right)\right) \stackrel{(a)}{=} \int_{0}^{1}\left|a_{1} \varepsilon_{1}(y)+a_{2} \varepsilon_{2}(y)\right|^{p} \mathrm{~d} y \stackrel{(b)}{\leq} \int_{0}^{1}\left(a_{1}\left|\varepsilon_{1}(y)\right|+\right. \\
\left.a_{2}\left|\varepsilon_{2}(y)\right|\right)^{p} \mathrm{~d} y \stackrel{(c)}{\leq} a_{1} \mathrm{~W}_{p}^{p}\left(\widetilde{q}, \mathcal{N}\left(\mu_{1}, \sigma_{1}^{2}\right)\right)+a_{2} \mathrm{~W}_{p}^{p}\left(\widetilde{q}, \mathcal{N}\left(\mu_{2}, \sigma_{2}^{2}\right)\right)
\end{gathered}
$$

where steps (a), (b) and (c) are obtained by applying Proposition 1, non-negativity
of the absolute value, and the convexity of $f(x)=x^{p}, p \geq 1$, over $\mathbb{R}_{+}$respectively. The equality at (b) holds iff $\varepsilon_{k}(y) \geq 0, k=1,2, \forall y \in[0,1]$, and (c)'s equality holds iff $\left|\varepsilon_{1}(y)\right|=\left|\varepsilon_{2}(y)\right|, \forall y \in[0,1]$. These two conditions for equality can't be attained simultaneously as otherwise it would contradict that $\mathcal{N}\left(\mu_{1}, \sigma_{1}^{2}\right)$ is different from $\mathcal{N}\left(\mu_{2}, \sigma_{2}^{2}\right)$. So, $\mathrm{W}_{p}^{p}(\widetilde{q}, \mathcal{N}), p \geq 1$, is strictly convex about $\mu$ and $\sigma$.

## C. 5 Proof of Variance Difference

Theorem 11. The variance of the Gaussian approximation to a univariate tilted distribution $\widetilde{q}(f)$ as estimated by $Q P$ and $E P$ satisfy $\sigma_{Q P}^{2} \leq \sigma_{E P}^{2}$.

Proof. Let $\mathcal{N}\left(\mu_{\mathrm{QP}}, \sigma_{\mathrm{QP}}^{2}\right)$ be the optimal Gaussian in QP. As per Proposition 1, we reformulate the $L_{2}$ WD based projection $\mathrm{W}_{2}^{2}\left(\widetilde{q}, \mathcal{N}\left(\mu_{\mathrm{QP}}, \sigma_{\mathrm{QP}}^{2}\right)\right)$ w.r.t. quantile functions:

$$
\begin{aligned}
\mathrm{W}_{2}^{2}\left(\widetilde{q}, \mathcal{N}\left(\mu_{\mathrm{QP}}, \sigma_{\mathrm{QP}}^{2}\right)\right)= & \int_{0}^{1}\left|F_{\tilde{q}}^{-1}(y)-\mu_{\mathrm{QP}}-\sqrt{2} \sigma_{\mathrm{QP}} \operatorname{erf}^{-1}(2 y-1)\right|^{2} \mathrm{~d} y \\
= & \int_{0}^{1} \underbrace{\left(F_{\widetilde{q}}^{-1}(y)-\mu_{\mathrm{QP}}\right)^{2}}_{\sigma_{\mathrm{EP}}^{2}}+\underbrace{\left(\sqrt{2} \sigma_{\mathrm{QP}} \operatorname{erf}^{-1}(2 y-1)\right)^{2}}_{\sigma_{\mathrm{QP}}^{2}} \\
& -\underbrace{2\left(F_{\tilde{q}}^{-1}(y)-\mu_{\mathrm{QP}}\right) \sqrt{2} \sigma_{\mathrm{QP}} \operatorname{erf}^{-1}(2 y-1)}_{(\mathrm{A})} \mathrm{d} y \\
= & \sigma_{\mathrm{EP}}^{2}-\sigma_{\mathrm{QP}}^{2},
\end{aligned}
$$

where for (A), we used $\int \mu_{\mathrm{QP}} \sigma_{\mathrm{QP}} \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y=0$ and the remaining factor can be easily shown to be equal to $2 \sigma_{\mathrm{QP}}^{2}$. Furthermore, due to the non-negativity of the WD , we have $\sigma_{\mathrm{EP}}^{2} \geq \sigma_{\mathrm{QP}}^{2}$, and the equality holds if and only if $\tilde{q}$ is Gaussian.

## C. 6 Proof of Locality Property

Theorem. Minimization of $\mathrm{W}_{2}^{2}(\widetilde{q}(f), \mathcal{N}(\boldsymbol{f}))$ w.r.t. $\mathcal{N}(\boldsymbol{f})$ results in $q^{\backslash i}\left(f_{\backslash i} \mid f_{i}\right)=$ $\mathcal{N}\left(f_{\backslash i} \mid f_{i}\right)$.

Proof. We first apply the decomposition of the $L_{2}$ norm to rewriting the $W_{2}^{2}(\widetilde{q}(f), \mathcal{N}(f))$ as below (see detailed derivations in Appendix C.6.2):

$$
\begin{equation*}
\mathrm{W}_{2}^{2}(\widetilde{q}, \mathcal{N})=\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\mathrm{W}_{2}^{2}\left(q_{i \mid i i^{\prime}}^{\backslash i} \mathcal{N}_{\langle i| i}\right)\right], \tag{С.3}
\end{equation*}
$$

where the prime indicates that the variable is from the Gaussian $\mathcal{N}$, and for simplification, we use the notation $\pi_{i}$ for the joint distribution $\pi\left(f_{i}, f_{i}^{\prime}\right)$ which belongs to a set of measures $U\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)$. Since $q^{\backslash i}(f)$ is known to be Gaussian, we define it in a
partitioned form:

$$
q^{\backslash i}(f) \equiv \mathcal{N}\left(\left.\left[\begin{array}{c}
f_{\backslash i}  \tag{C.4}\\
f_{i}
\end{array}\right] \right\rvert\,\left[\begin{array}{c}
m_{\backslash i} \\
m_{i}
\end{array}\right],\left[\begin{array}{cc}
S_{\backslash i} & S_{\backslash i i} \\
S_{\backslash i i}^{\top} & S_{i}
\end{array}\right]\right)
$$

and the conditional $q^{\backslash i}\left(f_{\backslash i} \mid f_{i}\right)$ is expressed as:

$$
\begin{array}{ll}
q^{\backslash i}\left(f_{\backslash \backslash \mid} \mid f_{i}\right)=\mathcal{N}\left(f_{\backslash i} \mid \boldsymbol{m}_{\backslash \backslash \mid i}, \boldsymbol{S}_{\backslash i \mid i}\right), & \boldsymbol{m}_{\backslash i \mid i}=\boldsymbol{m}_{\backslash i}+\boldsymbol{S}_{\backslash i i} S_{i}^{-1}\left(f_{i}-m_{i}\right) \equiv \boldsymbol{a} f_{i}+\boldsymbol{b},  \tag{C.5}\\
& \boldsymbol{S}_{\backslash i \mid i}=\boldsymbol{S}_{\backslash i}-\boldsymbol{S}_{\backslash i i} S_{i}^{-1} \boldsymbol{S}_{\backslash i i}^{\mathrm{T}} .
\end{array}
$$

We define a similar partitioned expression for the Gaussian $\mathcal{N}\left(f^{\prime}\right)$ by adding primes to variables and parameters on the r.h.s. of Equation (C.4), and as a result, the conditional $\mathcal{N}\left(f_{\backslash i}^{\prime} \mid f_{i}^{\prime}\right)$ is written as:

$$
\begin{array}{cl}
\mathcal{N}\left(\boldsymbol{f}_{\backslash \backslash \mid}^{\prime} \mid f_{i}^{\prime}\right)=\mathcal{N}\left(\boldsymbol{m}_{\backslash \backslash \mid i}^{\prime}, \boldsymbol{S}_{\backslash i \mid i}^{\prime}\right), & \boldsymbol{m}_{\backslash i \mid i}^{\prime}=\boldsymbol{m}_{\backslash i}^{\prime}+\boldsymbol{S}_{\backslash i i}^{\prime} S_{i}^{\prime-1}\left(f_{i}^{\prime}-m_{i}^{\prime}\right) \equiv \boldsymbol{a}^{\prime} f_{i}^{\prime}+\boldsymbol{b}^{\prime}, \\
\boldsymbol{S}_{\backslash i \mid i}^{\prime}=\boldsymbol{S}_{\backslash i}^{\prime}-\boldsymbol{S}_{\backslash i i}^{\prime} S_{i}^{\prime-1} \boldsymbol{S}_{\backslash i i}^{\prime T} . \tag{C.7}
\end{array}
$$

Given the above definitions, we exploit Proposition 2 to take the means out of the $L_{2}$ WD on the r.h.s. of Equation (C.3):

$$
\mathrm{W}_{2}^{2}(\widetilde{q}, \mathcal{N})=\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\left\|\boldsymbol{m}_{\backslash i \mid i}-\boldsymbol{m}_{\backslash i \mid i}^{\prime}\right\|_{2}^{2}\right]+\underbrace{\mathrm{W}_{2}^{2}\left(\mathcal{N}\left(\mathbf{0}, S_{\backslash i \mid i}\right), \mathcal{N}\left(\mathbf{0}, S_{\backslash i \mid i}^{\prime}\right)(\mathrm{C} .8)\right.}_{(\mathrm{A})}
$$

Minimizing this function requires optimizing $m_{i}^{\prime}, \boldsymbol{m}_{\backslash{ }_{i}}^{\prime}, S_{i}^{\prime}, \boldsymbol{S}_{\backslash i}^{\prime}$ and $\boldsymbol{S}_{\backslash i i}^{\prime}$. As $\boldsymbol{S}_{\backslash i}^{\prime}$ is only contained in $S_{\backslash i \mid i}$ and isolated into the term (A), it can be optimized by simply setting

$$
\begin{equation*}
S_{\backslash i \mid i}^{\prime}=S_{\backslash i \mid i} \stackrel{\text { Equation }}{\Longrightarrow}{ }^{(\mathrm{C} .7)} \boldsymbol{S}_{\backslash i}^{(n) *}=S_{\backslash i \mid i}+S_{\backslash i}^{\prime} S_{i}^{\prime-1} \boldsymbol{S}_{\backslash i i}^{\prime \text { T }} . \tag{C.9}
\end{equation*}
$$

As a result, (A) is minimized to zero. Next, we plug in expressions of $\boldsymbol{m}_{\backslash i \mid i}$ and $\boldsymbol{m}_{\backslash i \mid i}^{\prime}$ (Equation (C.5) and Equation (C.6)) into optimized Equation (C.8):

$$
\begin{equation*}
\min _{S_{\backslash_{i}^{\prime}}}(\mathrm{C} .8)=\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\left\|\boldsymbol{a} f_{i}-\boldsymbol{a}^{\prime} f_{i}^{\prime}+\boldsymbol{b}-\boldsymbol{b}^{\prime}\right\|_{2}^{2}\right] \tag{C.10}
\end{equation*}
$$

where $\boldsymbol{m}_{\backslash i}^{\prime}$ is only contained by $\boldsymbol{b}^{\prime}$. Thus, we can optimize it by zeroing the derivative of the above function about $\boldsymbol{m}_{\backslash i}^{\prime}$, which results in:

$$
\begin{equation*}
\boldsymbol{b}^{\prime}=\boldsymbol{b}+\boldsymbol{a} \mu_{\widetilde{q}_{i}}-\boldsymbol{a}^{\prime} m_{i}^{\prime} \stackrel{\text { Equation (C.6) }}{\Longrightarrow} \boldsymbol{m}_{\backslash i}^{(n) *}=\boldsymbol{S}_{\backslash i i}^{\prime} S_{i}^{\prime-1} m_{i}^{\prime}+\boldsymbol{b}+\boldsymbol{a} \mu_{\widetilde{q}_{i}}-\boldsymbol{a}^{\prime} m_{i}^{\prime} \tag{C.11}
\end{equation*}
$$

where $\mu_{\widetilde{q}_{i}}$ is the mean of $\widetilde{q}\left(f_{i}\right)$. The minimum value of Equation (C.10) thereby is (see
details in subsection C.6.3):

$$
\begin{equation*}
\min _{\boldsymbol{m}_{i}^{\prime}}(\mathrm{C} .10)=\left(1+\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\right) \mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}-\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\left[\sigma_{\widetilde{q}_{i}}^{2}+S_{i}^{\prime}+\left(\mu_{\widetilde{q}_{i}}-m_{i}^{\prime}\right)^{2}\right] \tag{C.12}
\end{equation*}
$$

where $\sigma_{\widetilde{q}_{i}}^{2}$ is the variance of $\widetilde{q}\left(f_{i}\right)$. This function can be further simplified using the quantile based reformulation of $W_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)$ (see details in Appendix C.6.4) which results in:

$$
\begin{equation*}
\text { (C.12) }=\mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\tilde{q}_{i}}^{2} \underbrace{-2^{\frac{3}{2}} \boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime} \boldsymbol{c}_{\tilde{q}_{i}} S_{i}^{\prime \frac{1}{2}}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}}_{(\mathrm{B})} . \tag{С.13}
\end{equation*}
$$

Now, we are left with optimizing $m_{i}^{\prime}, S_{i}^{\prime}$ and $S_{\langle i i}^{\prime}$. To optimize $S_{\backslash i i}^{\prime}$, which only exists in the above term (B), we zero the derivative of (B) w.r.t. $S_{{ }_{i i}}^{\prime}$ and this yields:

$$
\begin{equation*}
\boldsymbol{a}^{\prime *}=2^{\frac{1}{2}}\left(S_{i}^{\prime}\right)^{-\frac{1}{2}} \mathcal{C}_{\tilde{q}_{i}} \boldsymbol{a} \stackrel{\text { Equation }(\mathrm{C} .6)}{\Longrightarrow} \boldsymbol{S}_{\backslash i i}^{\prime *}=\left(2 S_{i}^{\prime}\right)^{\frac{1}{2}} \tilde{\mathcal{q}}_{\tilde{q}_{i}} \boldsymbol{a}, \tag{C.14}
\end{equation*}
$$

and the minimum value of Equation (C.13) is

$$
\begin{equation*}
\min _{S_{\backslash i i}^{\prime}}(\mathrm{C} .13)=\mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)+\|\boldsymbol{a}\|_{2}^{2}\left(\sigma_{\tilde{q}_{i}}^{2}-2 c_{\widetilde{q}_{i}}^{2}\right) . \tag{C.15}
\end{equation*}
$$

The results of optimizing $m_{i}^{\prime}$ and $S_{i}^{\prime}$ in the above equation have already been provided in Equation (5.1): $m_{i}^{* *}=\mu_{\widetilde{q}_{i}}$ and $S_{i}^{\prime *}=2 c_{\widetilde{q}_{i}}^{2}$. By plugging them into Equation (C.14) and Equation (C.11), we have $\boldsymbol{a}^{* *}=\boldsymbol{a}$ and $\boldsymbol{b}^{\boldsymbol{*}}=\boldsymbol{b}$. Finally, using Equation (C.9), we obtain $q^{\backslash i}\left(f_{\backslash i} \mid f_{i}\right)=\mathcal{N}\left(f_{\backslash i} \mid \boldsymbol{a} f_{i}+\boldsymbol{b}, \boldsymbol{S}_{\backslash i \mid}\right)=\mathcal{N}\left(\boldsymbol{f}_{\backslash i} \mid \boldsymbol{a}^{\prime} f_{i}+\boldsymbol{b}^{\prime}, \boldsymbol{S}_{\backslash i \mid i}^{\prime}\right)=\mathcal{N}\left(f_{\backslash i} \mid f_{i}\right)$, which concludes the proof.

## C.6.1 Details of Equation 5.2

$$
\begin{align*}
& \operatorname{KL}(\widetilde{q}(f) \| \mathcal{N}(f))=\int \widetilde{q}(f) \log \frac{\widetilde{q}\left(f_{\backslash i} \mid f_{i}\right) \widetilde{q}\left(f_{i}\right)}{\mathcal{N}\left(f_{\backslash i} \mid f_{i}\right) \mathcal{N}\left(f_{i}\right)} \mathrm{d} f \\
& =\int \widetilde{\mathfrak{q}}\left(f_{i}\right) \log \frac{\widetilde{q}\left(f_{i}\right)}{\mathcal{N}\left(f_{i}\right)} \mathrm{d} f_{i}+\int \widetilde{q}\left(f_{i}\right) \int \widetilde{q}\left(f_{\backslash i} \mid f_{i}\right) \log \frac{\widetilde{q}\left(f_{\backslash i} \mid f_{i}\right)}{\mathcal{N}\left(f_{\backslash \backslash i} \mid f_{i}\right)} \mathrm{d} f_{\backslash i} \mathrm{~d} f_{i} \\
& =\operatorname{KL}\left(\widetilde{q}\left(f_{i}\right) \| \mathcal{N}\left(f_{i}\right)\right)+\mathbb{E}_{\widetilde{q}\left(f_{i}\right)}\left[\operatorname{KL}\left(\widetilde{q}\left(f_{\backslash i} \mid f_{i}\right) \| \mathcal{N}\left(f_{\backslash i} \mid f_{i}\right)\right)\right] \\
& \widetilde{q}\left(f_{\backslash i} \mid f_{i}\right)=\frac{\widetilde{q}(f)}{\widetilde{q}\left(f_{i}\right)} \propto \frac{p(f) p\left(y_{i} \mid f_{i}\right) \prod_{j \neq i} t_{j}(f)}{\left.q^{\backslash}\left(f_{i}\right) p\left(y_{i}\right\rangle f_{i}\right)} \\
& =q^{\backslash i}\left(f_{\backslash i} \mid f_{i}\right) . \tag{C.16}
\end{align*}
$$

## C.6.2 Details of Equation C. 3

$$
\begin{aligned}
\mathrm{W}_{2}^{2}(\widetilde{q}(f), \mathcal{N}(f)) & \equiv \inf _{\pi \in U(\tilde{q}, \mathcal{N})} \mathbb{E}_{\pi}\left(\left\|f-f^{\prime}\right\|_{2}^{2}\right) \\
& =\inf _{\pi \in U(\tilde{q}, \mathcal{N})} \mathbb{E}_{\pi}\left(\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}\right)+\mathbb{E}_{\pi}\left(\left\|f_{\backslash i}-f_{\backslash i}^{\prime}\right\|_{2}^{2}\right) \\
& \stackrel{(a)}{=} \inf _{\pi \in U(\tilde{q}, \mathcal{N})} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\mathbb{E}_{\pi_{\backslash|i|}}\left(\left\|f_{\backslash i}-f_{\backslash i}^{\prime}\right\|_{2}^{2}\right)\right] \\
& \stackrel{(\mathrm{b})}{=} \inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\inf _{\pi_{\backslash \mid i}} \mathbb{E}_{\pi_{\backslash \mid i}}\left(\left\|f_{\backslash i}-f_{\backslash i}^{\prime}\right\|_{2}^{2}\right)\right] \\
& =\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\mathrm{W}_{2}^{2}\left(\widetilde{q}_{\backslash|i| i}, \mathcal{N}_{\backslash|i|}\right)\right] \\
& \stackrel{(\mathrm{c})}{=} \inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\mathrm{W}_{2}^{2}\left(q_{\backslash i \mid i}^{\backslash i}, \mathcal{N}_{\backslash i \mid i}\right)\right],
\end{aligned}
$$

where the superscript prime indicates that the variable is from the Gaussian $\mathcal{N}$. In (a), $\pi_{i}=\pi\left(f_{i}, f_{i}^{\prime}\right)$ and $\pi_{i \mid i}=\pi\left(f_{\backslash i}, f_{i i}^{\prime} \mid f_{i}, f_{i}^{\prime}\right)$. In (b), the first and the second inf are over $U\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)$ and $U\left(\widetilde{q}_{\backslash i \mid i}, \mathcal{N}_{\backslash i \mid i}\right)$ respectively. (c) is due to $\widetilde{q}\left(f_{\backslash i} \mid f_{i}\right)$ being equal to $q^{\backslash i}\left(f_{\backslash i} \mid f_{i}\right)$ (refer to Equation (C.16)).

## C.6.3 Details of Equation C. 12

$\min _{\boldsymbol{m}_{i i}}$ Equation (C.10)

$$
\begin{aligned}
& =\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}+\left\|\boldsymbol{a}\left(f_{i}-\mu_{\widetilde{q}_{i}}\right)-\boldsymbol{a}^{\prime}\left(f_{i}^{\prime}-m_{i}^{\prime}\right)\right\|_{2}^{2}\right] \\
& =\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}\right]+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}-2 \boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime} \mathbb{E}_{\pi_{i}}\left(f_{i} f_{i}^{\prime}-\mu_{\widetilde{q}_{i}} m_{i}^{\prime}\right) \\
& =\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}\right]+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}+\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime} \mathbb{E}_{\pi_{i}}\left(\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}-f_{i}^{2}-\left(f_{i}^{\prime}\right)^{2}+2 \mu_{\widetilde{q}_{i}} m_{i}^{\prime}\right) \\
& =\inf _{\pi_{i}} \mathbb{E}_{\pi_{i}}\left[\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}\right]+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}+\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime} \mathbb{E}_{\pi_{i}}\left(\left\|f_{i}-f_{i}^{\prime}\right\|_{2}^{2}-\left(f_{i}-\mu_{\widetilde{q}_{i}}^{2}\right)^{2}-\right. \\
& \\
& \left.2 f_{i} \mu_{\widetilde{q}_{i}}+\mu_{\widetilde{q}_{i}}^{2}-\left(f_{i}^{\prime}-m_{i}^{\prime}\right)^{2}-2 f_{i}^{\prime} m_{i}^{\prime}+\left(m_{i}^{\prime}\right)^{2}+2 \mu_{\widetilde{q}_{i}} m_{i}^{\prime}\right) \\
& =\left(1+\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\right) \mathrm{W}_{2}^{2}\left(\widetilde{q_{i}}, \mathcal{N}_{i}\right)+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}-\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\left(\sigma_{\widetilde{q}_{i}}^{2}+\mu_{\widetilde{q}_{i}}^{2}+S_{i}^{\prime}+\left(m_{i}^{\prime}\right)^{2}-2 \mu_{\widetilde{q}_{i}} m_{i}^{\prime}\right) \\
& =\left(1+\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\right) \mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}-\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\left[\sigma_{\widetilde{q}_{i}}^{2}+S_{i}^{\prime}+\left(\mu_{\widetilde{q}_{i}}-m_{i}^{\prime}\right)^{2}\right]
\end{aligned}
$$

## C.6.4 Details of Equation C. 12

We first use Proposition 1 to reformulate the $L_{2} W D W_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)$ as:

$$
\begin{aligned}
\mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right) & =\int_{0}^{1}\left(F_{\widetilde{q}_{i}}^{-1}(y)-m_{i}^{\prime}-\sqrt{2 S_{i}^{\prime}} \operatorname{erf}^{-1}(2 y-1)\right)^{2} \mathrm{~d} y, \\
& =\int_{0}^{1}\left(F_{\widetilde{q}_{i}}^{-1}(y)-m_{i}^{\prime}\right)^{2}+2 S_{i}^{\prime} \operatorname{erf}^{-1}(2 y-1)^{2}-2 \sqrt{2 S_{i}^{\prime}} \operatorname{erf}^{-1}(2 y-1)\left(F_{\widetilde{q}_{i}}^{-1}(y)-m_{i}^{\prime}\right) \mathrm{d} y, \\
& =\int_{0}^{1}\left(F_{\widetilde{q}_{i}}^{-1}(y)-\mu_{\widetilde{q}_{i}}+\mu_{\widetilde{q}_{i}}-m_{i}^{\prime}\right)^{2} \mathrm{~d} y+S_{i}^{\prime}-2 \sqrt{2 S_{i}^{\prime}} c_{\widetilde{q}_{i}}^{\prime} \\
& =\sigma_{\widetilde{q}_{i}}^{2}+\left(\mu_{\widetilde{q}_{i}}-m_{i}^{\prime}\right)^{2}+S_{i}^{\prime}-2 c_{\widetilde{q}_{i}} \sqrt{2 S_{i}^{\prime}}
\end{aligned}
$$

where $F_{\widetilde{q}_{i}}^{-1}(y)$ is the quantile function of $\widetilde{q}\left(f_{i}\right)$ and $c_{\widetilde{q}_{i}} \equiv \int_{0}^{1} F_{\widetilde{q}_{i}}^{-1}(y) \operatorname{erf}^{-1}(2 y-1) \mathrm{d} y$. Next, we plug this reformulation into Equation (C.12):

Equation (C.12) $=\mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)+\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime} \mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}-\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\left[\sigma_{\widetilde{q}_{i}}^{2}+S_{i}^{\prime}+\left(\mu_{\widetilde{q}_{i}}-m_{i}^{\prime}\right)^{2}\right]$

$$
\begin{aligned}
&= \mathrm{W}_{2}^{2}\left(\widetilde{q}_{i}, \mathcal{N}_{i}\right)+\boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}\left[\sigma_{\tilde{q}_{i}}^{2}+\left(\mu_{\tilde{q}_{i}}-m_{i}^{\prime}\right)^{2}+S_{i}^{\prime}\right. \\
&\left.-2 c_{\widetilde{q}_{i}} \sqrt{2 S_{i}^{\prime}}\right]+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\tilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime} \\
&=\left.\mathrm{W}_{2}^{2}\left(\widetilde{q}_{\tilde{q}_{i}}^{2}, \mathcal{N}_{i}\right)-2 S_{i}^{\prime}+\left(\tilde{q}_{\tilde{q}_{i}} \sqrt{2 \overparen{q}_{i}}-m_{i}^{\prime}\right)^{2}\right] \\
& \boldsymbol{a}^{\mathrm{T}} \boldsymbol{a}^{\prime}+\|\boldsymbol{a}\|_{2}^{2} \sigma_{\widetilde{q}_{i}}^{2}+\left\|\boldsymbol{a}^{\prime}\right\|_{2}^{2} S_{i}^{\prime}
\end{aligned}
$$

## C. 7 More Details of EP

We use the expressions $\widetilde{q}(\boldsymbol{f})=q^{\backslash i}(\boldsymbol{f}) p\left(y_{i} \mid f_{i}\right) / Z_{\widetilde{q}}$ and $q^{\backslash i}(\boldsymbol{f})=q(\boldsymbol{f}) /\left(t_{i}\left(f_{i}\right) Z_{q^{i}}\right)$, and the derivation of $\operatorname{KL}(\widetilde{q}(f) \| q(f))=\operatorname{KL}\left(\widetilde{q}\left(f_{i}\right) \| q\left(f_{i}\right)\right)$ is shown as below:

$$
\begin{aligned}
\operatorname{KL}(\widetilde{q}(f) \| q(f)) & =\int \widetilde{q}(f) \log \frac{q^{\backslash i}(f) p\left(y_{i} \mid f_{i}\right)}{Z_{\tilde{q} q} q(f)} \mathrm{d} f \\
& =\int \widetilde{q}(f) \log \frac{q(f) p\left(y_{i} \mid f_{i}\right)}{Z_{q q^{i}} Z_{\tilde{q} q} \tilde{}(f) t_{i}\left(f_{i}\right)} \mathrm{d} f \\
& =\int \widetilde{q}\left(f_{i}\right) \log \frac{p\left(y_{i} \mid f_{i}\right)}{Z_{q^{\backslash i}} Z_{\tilde{q}} t_{i}\left(f_{i}\right)} \mathrm{d} f_{i} \\
& =\int \widetilde{q}\left(f_{i}\right) \log \frac{q^{\backslash i}\left(f_{i}\right) p\left(y_{i} \mid f_{i}\right)}{Z_{q^{\backslash i}} Z_{\tilde{q} q} \backslash i\left(f_{i}\right) t_{i}\left(f_{i}\right)} \mathrm{d} f_{i} \\
& =\int \widetilde{q}\left(f_{i}\right) \log \frac{\widetilde{q}\left(f_{i}\right)}{q\left(f_{i}\right)} \mathrm{d} f_{i} \\
& =\operatorname{KL}\left(\widetilde{q}\left(f_{i}\right) \| q\left(f_{i}\right)\right)
\end{aligned}
$$

## C. 8 Predictive Distributions of Poisson Regression

Given the approximate predictive distribution $f\left(x_{*}\right)=\mathcal{N}\left(\mu_{*}, \sigma_{*}^{2}\right)$ and the relation $g(f)=f^{2}$, it is straightforward to derive the corresponding $g\left(\boldsymbol{x}_{*}\right) \sim \operatorname{Gamma}\left(k_{*}, c_{*}\right)^{1}$ where the shape $k_{*}$ and the scale $c_{*}$ are expressed as [Walder and Bishop, 2017; Zhang et al., 2020b]:

$$
k_{*}=\frac{\left(\mu_{*}^{2}+\sigma_{*}^{2}\right)^{2}}{2 \sigma_{*}^{2}\left(2 \mu_{*}^{2}+\sigma_{*}^{2}\right)}, \quad c_{*}=\frac{2 \sigma_{*}^{2}\left(2 \mu_{*}^{2}+\sigma_{*}^{2}\right)}{\mu_{*}^{2}+\sigma_{*}^{2}} .
$$

Furthermore, the predictive distribution of the count value $y \in \mathbb{N}$ can also be derived straightforwardly:

$$
\begin{aligned}
p(y) & =\int_{0}^{\infty} p\left(g_{*}\right) p\left(y \mid g_{*}\right) \mathrm{d} g_{*} \\
& =\int \operatorname{Gamma}\left(g_{*} \mid k_{*}, c_{*}\right) \operatorname{Poisson}\left(y \mid g_{*}\right) \mathrm{d} g_{*} \\
& =\frac{c_{*}^{y}\left(c_{*}+1\right)^{-k_{*}-y} \Gamma\left(k_{*}+y\right)}{y!\Gamma\left(k_{*}\right)}=\mathrm{NB}\left(y \mid k_{*}, c_{*} /\left(1+c_{*}\right)\right),
\end{aligned}
$$

where $g_{*}=g\left(\boldsymbol{x}_{*}\right)$ and NB denotes the negative binomial distribution. The mode is obtained as $\left\lfloor c_{*}\left(k_{*}-1\right)\right\rfloor$ if $k_{*}>1$ else 0 .

## C. 9 Proof of Corollary 3.2

Since the site approximations of both EP and QP are Gaussian, we may analyse the predictive variances using results from the regression with Gaussian likelihood function case, namely the well known Equation (3.61) in [Rasmussen and Williams, 2005]:

$$
\begin{equation*}
\sigma^{2}\left(f_{*}\right)=k\left(\boldsymbol{x}_{*}, \boldsymbol{x}_{*}\right)-\boldsymbol{k}_{*}^{\mathrm{T}}(K+\widetilde{\Sigma})^{-1} \boldsymbol{k}_{*}, \tag{C.17}
\end{equation*}
$$

where $f_{*}=f\left(\boldsymbol{x}_{*}\right)$ is the evaluation of the latent function at $\boldsymbol{x}_{*}$ and $\boldsymbol{k}_{*}=\left[k\left(\boldsymbol{x}_{*}, \boldsymbol{x}_{i}\right)\right]_{i=1}^{N^{\mathrm{T}}}$ is the covariance vector between the test data $\boldsymbol{x}_{*}$ and the training data $\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{N}, K$ is the prior covariance matrix and $\widetilde{\Sigma}$ is the diagonal matrix with elements of site variances $\widetilde{\sigma}_{i}^{2}$.

After updating the parameters of a site function $t_{i}\left(f_{i}\right)$, the term $(K+\widetilde{\Sigma})^{-1}$ is updated to $\left(K+\widetilde{\Sigma}+\left(\widetilde{\sigma}_{i, \text { new }}^{2}-\widetilde{\sigma}_{i}^{2}\right) \boldsymbol{e}_{i} \boldsymbol{e}_{i}^{T}\right)^{-1}$ where $\widetilde{\sigma}_{\text {inew }}$ is the site variance estimated by EP or QP and $e_{i}$ is a unit vector in direction $i$. Using the Woodbury, Sherman \& Morrison formula [Rasmussen and Williams, 2005, A.9], we rewrite ( $K+\widetilde{\Sigma}+\left(\widetilde{\sigma}_{i, \text { new }}^{2}-\right.$ $\left.\left.\widetilde{\sigma}_{i}^{2}\right) \boldsymbol{e}_{i} e_{i}^{\mathrm{T}}\right)^{-1}$ as

$$
\begin{aligned}
\left(K+\widetilde{\Sigma}+\left(\widetilde{\sigma}_{i \text {,new }}^{2}-\widetilde{\sigma}_{i}^{2}\right) \boldsymbol{e}_{i} \boldsymbol{e}_{i}^{\mathrm{T}}\right)^{-1} \equiv & \left(A^{-1}+\left(\widetilde{\sigma}_{i, \text { new }}^{2}-\widetilde{\sigma}_{i}^{2}\right) \boldsymbol{e}_{i} \boldsymbol{e}_{i}^{\mathrm{T}}\right)^{-1} \\
& =A-A \boldsymbol{e}_{i}\left[\left(\widetilde{\sigma}_{i, \text { new }}^{2}-\widetilde{\sigma}_{i}^{2}\right)^{-1}+\boldsymbol{e}_{i}^{\mathrm{T}} A \boldsymbol{e}_{i}\right]^{-1} \boldsymbol{e}_{i}^{\mathrm{T}} A
\end{aligned}
$$

[^3]\[

$$
\begin{aligned}
\equiv & A-\boldsymbol{s}_{i}\left[\left(\widetilde{\sigma}_{i, \text { new }}^{2}-\widetilde{\sigma}_{i}^{2}\right)^{-1}+A_{i i}\right]^{-1} \boldsymbol{s}_{i}^{\mathrm{T}} \\
& =A-\frac{1}{\left(\widetilde{\sigma}_{i, \text { new }}^{2}-\widetilde{\sigma}_{i}^{2}\right)^{-1}+A_{i i}} \boldsymbol{s}_{i} \boldsymbol{s}_{i}^{\mathrm{T}}
\end{aligned}
$$
\]

where $A=(K+\widetilde{\Sigma})^{-1}$ and $s_{i}$ is the $i^{\prime}$ th column of $A$. Putting the above expression into Equation (C.17), we have that the predictive variance is updated according to:

$$
\sigma_{\text {new }}^{2}\left(f_{*}\right)=k\left(\boldsymbol{x}_{*}, \boldsymbol{x}_{*}\right)-\boldsymbol{k}_{*}^{\mathrm{T}} A \boldsymbol{k}_{*}+\frac{1}{\left(\widetilde{\sigma}_{i, \text { new }}^{2}-\widetilde{\sigma}_{i}^{2}\right)^{-1}+A_{i i}} \boldsymbol{k}_{*}^{\mathrm{T}} \boldsymbol{s}_{i} \boldsymbol{s}_{i}^{\mathrm{T}} \boldsymbol{k}_{*}
$$

In EP and QP, the first two terms on the r.h.s. of the above equation are equivalent. As the site variance provided by QP is less or equal to that by EP, i.e., , $\widetilde{\sigma}_{i, \mathrm{QP}}^{2} \leq \widetilde{\sigma}_{i, \mathrm{EP}}^{2}$, the third term on the r.h.s. for QP is less or equal to that for EP. Therefore, the predictive variance of QP is less or equal to that of EP: $\sigma_{\mathrm{QP}}^{2}\left(f_{*}\right) \leq \sigma_{\mathrm{EP}}^{2}\left(f_{*}\right)$.

## C. 10 Lookup Tables

To speed up updating variances $\sigma_{\mathrm{QP}}^{2}$ in QP, we pre-compute the integration in Equation (5.1) over a grid of cavity parameters $\mu$ and $\sigma$, and store the results into lookup tables. Consequently, each update step obtains $\sigma_{\mathrm{QP}}^{2}$ simply based on the lookup tables. Concretely, for the GP binary classification, we compute Equation (5.1) with $\mu, \sigma$ and $y$ varying from -10 to $10,0.1$ to 10 and $\{-1,1\}$ respectively. $\mu$ and $\sigma$ vary in a linear scale and a $\log 10$ scale respectively, and both have a step size of 0.001 . The resulting lookup tables has a size of $20001 \times$ 2001. In a similar way, we make the lookup table for the Poisson regression. In the experiments, we exploit the linear interpolation to fit $\sigma_{\mathrm{QP}}^{2}$ given $\mu \in[-10,10]$ and $\sigma \in[0.1,10]$, and if $\mu$ and $\sigma$ lie out of the lookup table, $\sigma_{\mathrm{QP}}^{2}$ is approximately computed by the EP update formula, i.e., $\sigma_{\mathrm{QP}}^{2} \approx \sigma_{\mathrm{EP}}^{2}$. On $\operatorname{Intel}(\mathrm{R})$ Xeon(R) CPU E5-2680 v4 @ 2.40 GHz , we observe the running time of EP and QP is almost the same.

```
Algorithm 2 Expectation (Quantile) Propagation
Input: \(p(\boldsymbol{f}), p\left(y_{i} \mid f_{i}\right), t_{i}\left(f_{i}\right), i=1, \cdots, N, \boldsymbol{\theta}\)
Output: \(q(f)\)
approximate posterior
    repeat
        compute \(q(\boldsymbol{f}) \propto p(\boldsymbol{f}) \prod_{i} t_{i}\left(f_{i}\right) \quad\) by (2.8)
        repeat
            for \(i=1\) to \(N\) do
            compute \(q^{\backslash i}\left(f_{i}\right) \propto q\left(f_{i}\right) / t_{i}\left(f_{i}\right) \quad\) cavity
            compute \(\widetilde{q}\left(f_{i}\right) \propto q^{\backslash i}\left(f_{i}\right) p\left(y_{i} \mid f_{i}\right) \quad\) tilted
            if \(E P\) then
                \(t_{i}\left(f_{i}\right) \propto \operatorname{proj}_{\mathrm{KL}}\left[\widetilde{q}\left(f_{i}\right)\right] / q^{\backslash i}\left(f_{i}\right) \quad\) by (2.10)(2.11)
            else if QP then
                \(t_{i}\left(f_{i}\right) \propto \operatorname{proj}_{\mathrm{W}}\left[\widetilde{q}\left(f_{i}\right)\right] / q^{\backslash i}\left(f_{i}\right) \quad\) by (5.1)(2.11)
            end if
            update \(q(f) \propto p(f) \prod_{i} t_{i}\left(f_{i}\right) \quad\) by (2.8)
        end for
        until convergence
        \(\boldsymbol{\theta}=\operatorname{argmax}_{\theta} \log q(\mathcal{D})\)
                            by (2.9)
    until convergence
    return \(q(f)\)
```


## Appendix: Kernel Maximum Moment Restriction

## D. 1 Integrally Strictly Positive Definite (ISPD) Kernels

Popular kernel functions that satisfy Assumption 1 are the Gaussian RBF kernel and Laplacian kernel

$$
k\left(z, z^{\prime}\right)=\exp \left(-\frac{\left\|z-z^{\prime}\right\|_{2}^{2}}{2 \sigma^{2}}\right), \quad k\left(z, z^{\prime}\right)=\exp \left(-\frac{\left\|z-z^{\prime}\right\|_{1}}{\sigma}\right)
$$

where $\sigma$ is a positive bandwidth parameter. Another important kernel is an inverse multiquadric (IMQ) kernel

$$
k\left(z, z^{\prime}\right)=\left(c^{2}+\left\|z-z^{\prime}\right\|_{2}^{2}\right)^{-\gamma}
$$

where $c$ and $\gamma$ are positive parameters [Steinwart and Christmann, 2008, Ch. 4]. This class of kernel functions is closely related to the notions of universal kernels [Steinwart, 2002] and characteristic kernels [Fukumizu et al., 2004]. The former ensures that kernel-based classification/regression algorithms can achieve the Bayes risk, whereas the latter ensures that the kernel mean embeddings can distinguish different probability measures. In principle, they guarantee that the corresponding RKHSs induced by these kernels are sufficiently rich for the tasks at hand. We refer the readers to Sriperumbudur et al. [2011] and Simon-Gabriel and Schölkopf [2018] for more details.

## D. 2 Detailed Proofs

This section contains detailed proofs of the results that are missing in the main paper. Most of the proofs on consistency and asymptotic normality take advantages of the useful resource by Newey and McFadden [1994]. Readers are referred to it for more detailed discussions on e.g. assumptions.

## D.2.1 Proof of Lemma 1

Proof. Since $\mathcal{H}_{k}$ is the RKHS, we can rewrite (6.2) as

$$
\begin{align*}
R_{k}(f) & =\sup _{h \in \mathcal{H}_{k},\|h\| \leq 1}\left(\mathbb{E}\left[(Y-f(X))\langle h, k(Z, \cdot)\rangle_{\mathcal{H}}\right]\right)^{2} \\
& =\sup _{h \in \mathcal{H}_{k},\|h\| \leq 1}\left(\langle h, \mathbb{E}[(Y-f(X)) k(Z, \cdot)]\rangle_{\mathcal{H}}\right)^{2} \\
& =\|\mathbb{E}[(Y-f(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2}, \tag{D.1}
\end{align*}
$$

where we used the reproducing property of $\mathcal{H}_{k}$ in the first equality and the fact that $\mathcal{H}_{k}$ is a vector space in the last equality. By assumption, $\mathbb{E}[(Y-f(X)) k(Z, \cdot)]$ is Bochner integrable [Steinwart and Christmann, 2008, Def. A.5.20]. Hence, we can write (D.1) as

$$
\begin{aligned}
\|\mathbb{E}[(Y-f(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2} & =\langle\mathbb{E}[(Y-f(X)) k(Z, \cdot)], \mathbb{E}[(Y-f(X)) k(Z, \cdot)]\rangle_{\mathcal{H}_{k}} \\
& =\mathbb{E}\left[\langle(Y-f(X)) k(Z, \cdot), \mathbb{E}[(Y-f(X)) k(Z, \cdot)]\rangle_{\mathcal{H}_{k}}\right] \\
& =\mathbb{E}\left[\left\langle(Y-f(X)) k(Z, \cdot),\left(Y^{\prime}-f\left(X^{\prime}\right)\right) k\left(Z^{\prime}, \cdot\right)\right\rangle_{\mathcal{H}_{k}}\right] \\
& =\mathbb{E}\left[(Y-f(X))\left(Y^{\prime}-f\left(X^{\prime}\right)\right) k\left(Z, Z^{\prime}\right)\right],
\end{aligned}
$$

as required.

## D.2.2 Proof of Theorem 6

Proof. First, the law of iterated expectation implies that
$\mathbb{E}[(Y-f(X)) k(Z, \cdot)]=\mathbb{E}_{Z}\left[\mathbb{E}_{X Y}[(Y-f(X)) k(Z, \cdot) \mid Z]\right]=\mathbb{E}_{Z}\left[\mathbb{E}_{X Y}[Y-f(X) \mid Z] k(Z, \cdot)\right]$.
By Lemma 1, we know that $R_{k}(f)=\|\mathbb{E}[(Y-f(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2}$. As a result, $R_{k}(f)=0$ if $\mathbb{E}[Y-f(X) \mid z]=0$ for $P_{Z}$-almost all $z$. To show the converse, we assume that $R_{k}(f)=0$ and rewrite it as

$$
R_{k}(f)=\iint_{\mathcal{Z}} g(z) k\left(z, z^{\prime}\right) g\left(z^{\prime}\right) \mathrm{d} z \mathrm{~d} z^{\prime}=0
$$

where we define $g(z):=\mathbb{E}_{X Y}[Y-f(X) \mid z] p(z)$. Since $k$ is ISPD by assumption, this implies that $g$ is a zero function with respect to $P_{Z}$, i.e., $\mathbb{E}[Y-f(X) \mid z]=0$ for $P_{\mathrm{Z}}$-almost all $z$.

## D.2.3 Convexity Result

Theorem 12. If $\mathcal{F}$ is a convex set and Assumptions 1,2 hold, then the risk $R_{k}$ given in (6.3) is strictly convex on $\mathcal{F}$.

Proof. Given $\alpha \in(0,1)$ and any functions $f, g: \mathcal{X} \rightarrow \mathbb{R}$, we will show that

$$
R_{k}(\alpha f+(1-\alpha) g)-\alpha R_{k}(f)-(1-\alpha) R_{k}(g)<0 .
$$

By Lemma 1, we know that $R_{k}(f)=\|\mathbb{E}[(Y-f(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2}$. Hence, we can rewrite the above function as

$$
\begin{aligned}
& R_{k}(\alpha f+(1-\alpha) g)-\alpha R_{k}(f)-(1-\alpha) R_{k}(g) \\
& =\|\mathbb{E}[(Y-\alpha f(X)-(1-\alpha) g(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2}-\alpha\|\mathbb{E}[(Y-f(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2} \\
& \quad-(1-\alpha)\|\mathbb{E}[(Y-g(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2} \\
& \stackrel{(\text { a) }}{=} \alpha(\alpha-1)\|\mathbb{E}[(Y-f(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2}+\alpha(\alpha-1)\|\mathbb{E}[(Y-g(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2} \\
& \quad-2 \alpha(\alpha-1)\langle\mathbb{E}[(Y-f(X)) k(Z, \cdot)], \mathbb{E}[(Y-g(X)) k(Z, \cdot)]\rangle_{\mathcal{H}_{k}} \\
& \stackrel{\text { (b) }}{=} \alpha(\alpha-1) \\
& \\
& \quad \underbrace{\|\mathbb{E}[(f(X)-g(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}^{2}}_{>0}<0
\end{aligned}
$$

The equality (a) is obtained by considering $Y=\alpha Y+(1-\alpha) Y$ in $\| \mathbb{E}[(Y-\alpha f(X)-$ $(1-\alpha) g(X)) k(Z, \cdot)] \|_{\mathcal{H}_{k}}^{2}$ on the left hand side of (a). We note that the right hand side of (a) is quadratic in $\|\mathbb{E}[(Y-f(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k}}$ and $\|\mathbb{E}[(Y-g(X)) k(Z, \cdot)]\|_{\mathcal{H}_{k^{\prime}}}$ and can be further expressed as a square binomial as the right hand side of $(b)$. Therefore, the convexity follows from the fact that $k$ is the ISPD kernel, $\|\mathbb{E}[(f(X)-g(X)) \mid Z]\|_{2} \neq 0$ and $\alpha(\alpha-1)<0$.

## D.2.4 Uniform Convergence of Risk Functionals

The results presented in this section are used to prove the consistency of $\hat{f}_{V}$ and $\hat{f}_{U}$.
Lemma 2 (Uniform consistency of $\widehat{R}_{V}(f)$ ). Assume that $\mathbb{E}\left[|Y|^{2}\right]<\infty, \mathcal{F}$ is compact, $\mathbb{E}\left[\sup _{f \in \mathcal{F}}|f(X)|^{2}\right]<\infty$, and Assumption 1 holds. Then, the risk $R_{k}(f)$ is continuous about $f \in \mathcal{F}$ and $\sup _{f \in \mathcal{F}}\left|\widehat{R}_{V}(f)-R_{k}(f)\right| \xrightarrow{\mathrm{p}} 0$.

Proof. First, let $u:=(x, y, z), u^{\prime}:=\left(x^{\prime}, y^{\prime}, z^{\prime}\right)$, and $h_{f}\left(u, u^{\prime}\right):=(y-f(x))\left(y^{\prime}-\right.$ $\left.f\left(x^{\prime}\right)\right) k\left(z, z^{\prime}\right)$ for some $(x, y, z),\left(x^{\prime}, y^{\prime}, z^{\prime}\right) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$. To prove that $\widehat{R}_{V}$ converges uniformly to $R_{k}$, we need to show that (i) $h_{f}\left(u, u^{\prime}\right)$ is continuous at each $f$ with probability one; (ii) $\mathbb{E}_{U, U^{\prime}}\left[\sup _{f \in \mathcal{F}}\left|h_{f}\left(U, U^{\prime}\right)\right|\right]<\infty$, and $\mathbb{E}_{U, U}\left[\sup _{f \in \mathcal{F}}\left|h_{f}(U, U)\right|\right]<\infty$ Newey and McFadden [1994, Lemma 8.5]. To this end, it is easy to see that

$$
\begin{aligned}
\left|h_{f}\left(u, u^{\prime}\right)\right| & =\left|(y-f(x))\left(y^{\prime}-f\left(x^{\prime}\right)\right) k\left(z, z^{\prime}\right)\right| \\
& \leq|y-f(x)|\left|y^{\prime}-f\left(x^{\prime}\right)\right|\left|k\left(z, z^{\prime}\right)\right| \\
& \leq|y-f(x)|\left|y^{\prime}-f\left(x^{\prime}\right)\right| \sqrt{k(z, z) k\left(z^{\prime}, z^{\prime}\right)} \\
& \leq(|y|+|f(x)|)\left(\left|y^{\prime}\right|+\left|f\left(x^{\prime}\right)\right|\right) \sqrt{k(z, z) k\left(z^{\prime}, z^{\prime}\right)} .
\end{aligned}
$$

The third inequality follows from the Cauchy-Schwarz inequality. Since $\mathcal{F}$ is compact, every $f \in \mathcal{F}$ has $f(x)$ bounded for $\|x\|<\infty$. In term of $k(\cdot, \cdot)$ is bounded as per Assumption 1, we have $h_{f}\left(u, u^{\prime}\right)<\infty$ and thus $h_{f}\left(u, u^{\prime}\right)$ continuous at each $f$ with
probability one. Furthermore, we obtain the following inequalities

$$
\begin{aligned}
& \mathbb{E}_{U, U^{\prime}}\left[\sup _{f \in \mathcal{F}}\left|h_{f}\left(U, U^{\prime}\right)\right|\right] \leq \leq \mathbb{E}\left[\sup _{f \in \mathcal{F}}(|Y|+|f(X)|)\left(\left|Y^{\prime}\right|+\left|f\left(X^{\prime}\right)\right|\right) \sqrt{k(Z, Z) k\left(Z^{\prime}, Z^{\prime}\right)}\right] \\
& \leq \mathbb{E}\left[\sup _{f \in \mathcal{F}}(|Y|+|f(X)|) \sup _{f \in \mathcal{F}}\left(\left|Y^{\prime}\right|+\left|f\left(X^{\prime}\right)\right|\right)\right] \sup _{z} k(z, z) \\
&=\mathbb{E}\left[\sup _{f \in \mathcal{F}}(|Y|+|f(X)|)\right]^{2} \sup _{z} k(z, z) \\
&=\mathbb{E}\left[|Y|+\sup _{f \in \mathcal{F}}|f(X)|\right]^{2} \sup _{z} k(z, z)<\infty \\
& \begin{aligned}
\mathbb{E}_{U, U}\left[\sup _{f \in \mathcal{F}}\left|h_{f}(U, U)\right|\right] \leq & \mathbb{E}_{u}\left[\sup _{f \in \mathcal{F}}(|Y|+|f(X)|)^{2}\right] \sup _{z} k(z, z) \\
& =\mathbb{E}\left[\left(|Y|+\sup _{f \in \mathcal{F}}|f(X)|\right)^{2}\right] \sup _{z} k(z, z) \\
& \leq 2\left(\mathbb{E}\left[|Y|^{2}\right]+\mathbb{E}\left[\sup _{f \in \mathcal{F}}|f(X)|^{2}\right]\right) \sup _{z} k(z, z)<\infty
\end{aligned}
\end{aligned}
$$

Hence, our assertion follows from Newey and McFadden [1994, Lemma 8.5].
Lemma 3 (Uniform consistency of $\widehat{R}_{U}(f)$ ). Assume that $\mathbb{E}[|Y|]<\infty, \mathcal{F}$ is compact, $\mathbb{E}[|f(X)|]<\infty$ and Assumption 1 holds. Then, $R_{k}(f)$ is continuous about $f$ and $\sup _{f \in \mathcal{F}} \mid \widehat{R}_{U}(f)-$ $R_{k}(f) \mid \xrightarrow{\mathrm{P}} 0$.

Proof. First, let $u:=(x, y, z), u^{\prime}:=\left(x^{\prime}, y^{\prime}, z^{\prime}\right)$, and $h_{f}\left(u, u^{\prime}\right):=(y-f(x))\left(y^{\prime}-\right.$ $\left.f\left(x^{\prime}\right)\right) k\left(z, z^{\prime}\right)$ for some $(x, y, z),\left(x^{\prime}, y^{\prime}, z^{\prime}\right) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$. To prove the uniform consistency of $\widehat{R}_{U}$, we need to show that (i) $h_{f}\left(u, u^{\prime}\right)$ is continuous at each $f$ with probability one; (ii) there is $d\left(u, u^{\prime}\right)$ with $\left|h_{f}\left(u, u^{\prime}\right)\right| \leq d\left(u, u^{\prime}\right)$ for all $f \in \mathcal{F}$ and $\mathbb{E}_{U, u^{\prime}}\left[d\left(U, U^{\prime}\right)\right]<\infty$ [Newey and McFadden, 1994, Lemma 2.4]; (iii) $\left(u_{i}, u_{j}\right)_{i \neq j}^{n, n}$ has strict stationarity and ergodicity in the sense of Newey and McFadden [1994, Footnote 18 in P.2129]. To this end, it is easy to see that

$$
\begin{aligned}
\left|h_{f}\left(u, u^{\prime}\right)\right| & =\left|(y-f(x))\left(y^{\prime}-f\left(x^{\prime}\right)\right) k\left(z, z^{\prime}\right)\right| \\
& \leq|y-f(x)|\left|y^{\prime}-f\left(x^{\prime}\right)\right|\left|k\left(z, z^{\prime}\right)\right| \\
& \leq|y-f(x)|\left|y^{\prime}-f\left(x^{\prime}\right)\right| \sqrt{k(z, z) k\left(z^{\prime}, z^{\prime}\right)} \\
& \leq(|y|+|f(x)|)\left(\left|y^{\prime}\right|+\left|f\left(x^{\prime}\right)\right|\right) \sqrt{k(z, z) k\left(z^{\prime}, z^{\prime}\right)} \equiv d\left(u, u^{\prime}\right) .
\end{aligned}
$$

The third inequality follows from the Cauchy-Schwarz inequality. Since $\mathcal{F}$ is compact, every $f \in \mathcal{F}$ has $f(x)$ bounded for $\|x\|<\infty$. In terms of $k(\cdot, \cdot)$ is bounded as per Assumption 1, we have $h_{f}\left(u, u^{\prime}\right)<\infty$ and thus it proves that (i) $h_{f}\left(u, u^{\prime}\right)$ is continuous
at each $f$ with probability one. To prove (ii) $\mathbb{E}_{U, U^{\prime}}\left[d\left(U, U^{\prime}\right)\right]<\infty$, we show that

$$
\mathbb{E}_{U, U^{\prime}}\left[d\left(U, U^{\prime}\right)\right] \leq \mathbb{E}[|Y|+|f(X)|]^{2} \sup _{z} k(z, z)<\infty
$$

Furthermore, we show that $\left(u_{i}, u_{j}\right)_{i \neq j}^{n, n}$ has strict stationarity and ergodicity. Strict stationarity means that the distribution of a set of data $\left(u_{i}, u_{j \neq i}\right)_{i=1, j=1}^{i+m, j+m^{\prime}}$ does not depend on the starting indices $i, j$ for any $m$ and $m^{\prime}$, which is easy to check. Ergodicity means that $\widehat{R}_{U}(f) \xrightarrow{\mathrm{p}} R_{k}(f)$ for all $f \in \mathcal{F}$ and $\mathbb{E}\left[\left|h_{f}\left(U, U^{\prime}\right)\right|\right]<\infty$. We have already shown that $h_{f}\left(u, u^{\prime}\right)$ is bounded and $\mathbb{E}\left[\left|h_{f}\left(U, U^{\prime}\right)\right|\right]<\infty$, so $\widehat{R}_{U}(f) \xrightarrow{\mathrm{p}} R_{k}(f)$ follows by Hoeffding [1963, P.25]. Therefore, ergodicity holds, and we have shown all conditions required by extended results of Newey and McFadden [1994, Lemma 2.4]. Then, it follows that $\sup _{f \in \mathcal{F}}\left|\widehat{R}_{U}(f)-R_{k}(f)\right| \xrightarrow{\mathrm{p}} 0$ and $R_{k}(f)$ is continuous.

## D.2.5 Indefiniteness of Weight Matrix $W_{U}$

Theorem 13. If Assumption 1 holds, $W_{U}$ is indefinite.
Proof. By definition, we have

$$
W_{U}=\frac{1}{n(n-1)}\left[K(z, z)-\operatorname{diag}\left(k\left(z_{1}, z_{1}\right), \ldots, k\left(z_{n}, z_{n}\right)\right)\right]=\frac{1}{n(n-1)} K_{U},
$$

where $\operatorname{diag}\left(a_{1}, \ldots, a_{n}\right)$ denotes an $n \times n$ diagonal matrix whose diagonal elements are $a_{1}, \ldots, a_{n}$. We can see that the diagonal elements of $K_{U}$ are zeros and therefore $\operatorname{trace}\left(W_{U}\right)=0$. Let us denote the eigenvalues of $W_{U}$ by $\left\{\lambda_{i}\right\}_{i=1}^{n}$. Since $\sum_{i=1}^{n} \lambda_{i}=$ trace $\left(W_{U}\right)$, we conclude that there exist both positive and negative eigenvalues (all eigenvalues being zeros yields trivial $W_{U}=\mathbf{0}$ ). As a result, $W_{U}$ is indefinite.

## D.2.6 Consistency of $\hat{f}_{V}$ with Convex $\Omega(f)$

Theorem 14 (Consistency of $\hat{f}_{V}$ with convex $\Omega(f)$ ). Assume that $\mathcal{F}$ is a convex set, $f^{*}$ is an interior point of $\mathcal{F}, \Omega(f)$ is convex about $f, \lambda \xrightarrow{\mathrm{p}} 0$ and Assumptions 1,2 holds. Then, $\hat{f}_{V}$ exists with probability approaching one and $\hat{f}_{V} \xrightarrow{p} f^{*}$.

Proof. Given $\Omega(f)$ is convex about $f$, we prove the consistency based on Newey and McFadden [1994, Theorem 2.7] which requires (i) $R_{k}(f)$ is uniquely maximized at $f^{*}$; (ii) $\widehat{R}_{V}(f)+\lambda \Omega(f)$ is convex; (iii) $\widehat{R}_{V}(f)+\lambda \Omega(f) \xrightarrow{\mathrm{p}} R_{k}(f)$ for all $f \in \mathcal{F}$.

Recall that $\widehat{R}_{V}(f)=\left\|\frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right) k\left(z_{i}, \cdot\right)\right\|_{\mathcal{H}_{k^{\prime}}}^{2}$ and by the law of large number, we have that $\frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right) k\left(z_{i}, \cdot\right) \xrightarrow{\mathrm{p}} \mathbb{E}[(Y-f(X)) k(Z, \cdot)]$. Then $\widehat{R}_{V}(f) \xrightarrow{\mathrm{p}} R_{k}(f)$ follows from the Continuous Mapping Theorem [Mann and Wald, 1943] based on the fact that the function $g(\cdot)=\|\cdot\|_{\mathcal{H}_{k}}^{2}$ is continuous. As $\lambda \xrightarrow{\mathrm{p}} 0$, we obtain (iii) $\widehat{R}_{V}(f)+\lambda \Omega(f) \xrightarrow{\mathrm{p}} R_{k}(f)$ by Slutsky's theorem [Van der Vaart, 2000, Lemma 2.8]. Besides, it is easy to see that $\widehat{R}_{V}(f)$ is convex because the weight matrix $W_{V}$ is positive definite, and (ii) $\widehat{R}_{V}(f)+\lambda \Omega(f)$ is convex due to convex $\Omega(f)$. Further, the condition
(i) directly follows from Theorem 12, and given that $f^{*}$ is an interior point of the convex set $\mathcal{F}$, our assertion follows from Newey and McFadden [1994, Theorem 2.7].

## D.2.7 Proof of Theorem 8

Proof. From the conditions of Lemma 2, we know that $\mathcal{F}$ is compact, $R_{k}(f)$ is continuous about $f$ and $\sup _{f \in \mathcal{F}}\left|\widehat{R}_{V}(f)-R_{k}(f)\right| \xrightarrow{\mathrm{p}} 0$. As Assumptions 1,2 hold, $R_{k}(f)$ is uniquely minimized at $f^{*}$. Based on the conditions that $\Omega(f)$ is bounded and $\lambda \xrightarrow{\mathrm{P}} 0$, we obtain by Slutsky's theorem that

$$
\sup _{f \in \mathcal{F}}\left|\widehat{R}_{V}(f)+\lambda \Omega(f)-R_{k}(f)\right| \leq \sup _{f \in \mathcal{F}}\left|\widehat{R}_{V}(f)-R_{k}(f)\right|+\lambda \sup _{f \in \mathcal{F}} \Omega(f) \xrightarrow{\mathrm{p}} 0 .
$$

Consequently, we assert the conclusion by Newey and McFadden [1994, Theorem 2.1].

## D.2.8 Consistency of $\hat{f}_{U}$

Theorem 15 (Consistency of $\hat{f}_{U}$ ). Assume that conditions of Lemma 3 and Assumption 2 hold, $\Omega(f)$ is a bounded function and $\lambda \xrightarrow{\mathrm{p}} 0$. Then $\hat{f}_{U} \xrightarrow{\mathrm{p}} f^{*}$.

Proof. By the conditions of Lemma 3, we know that $\mathcal{F}$ is compact, $R_{k}(f)$ is continuous about $f$ and $\sup _{f \in \mathcal{F}}\left|\widehat{R}_{U}(f)-R_{k}(f)\right| \xrightarrow{p} 0$. As Assumptions 1, 2 hold, $R_{k}(f)$ is uniquely minimized at $f^{*}$. Based on the conditions that $\Omega(f)$ is bounded and $\lambda \xrightarrow{\mathrm{p}} 0$, we obtain by Slutsky's theorem that

$$
\sup _{f \in \mathcal{F}}\left|\widehat{R}_{U}(f)+\lambda \Omega(f)-R_{k}(f)\right| \leq \sup _{f \in \mathcal{F}}\left|\widehat{R}_{U}(f)-R_{k}(f)\right|+\lambda \sup _{f \in \mathcal{F}} \Omega(f) \xrightarrow{\mathrm{p}} 0 .
$$

Consequently, we assert the conclusion by Newey and McFadden [1994, Theorem 2.1].

## D.2.9 Asymptotic Normality of $\hat{\theta}_{U}$

In this section, we consider the regularized U -statistic risk $\widehat{R}_{U, \lambda}\left(f_{\theta}\right)$. For $u_{i}:=\left(x_{i}, y_{i}, z_{i}\right)$ and $u_{j}:=\left(x_{j}, y_{j}, z_{j}\right)$, we express it in a compact form

$$
\begin{aligned}
& \widehat{R}_{U, \lambda}\left(f_{\theta}\right):=\underbrace{\frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j \neq i}^{n} h_{\theta}\left(u_{i}, u_{j}\right)}_{\widehat{R}_{u}\left(f_{\theta}\right)}+\lambda \Omega(\theta) \\
& h_{\theta}\left(u_{i}, u_{j}\right):=\left(y_{i}-f_{\theta}\left(x_{i}\right)\right) k\left(z_{i}, z_{j}\right)\left(y_{j}-f_{\theta}\left(x_{j}\right)\right) .
\end{aligned}
$$

We will assume that $f_{\theta}$ and $\Omega(\theta)$ are twice continuously differentiable about $\theta$. The first-order derivative $\nabla_{\theta} \widehat{R}_{U}\left(f_{\theta}\right)$ can also be written as

$$
\begin{aligned}
\nabla_{\theta} \widehat{R}_{U, \lambda}\left(f_{\theta}\right) & =\underbrace{\frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j \neq i}^{n} \nabla_{\theta} h_{\theta}\left(u_{i}, u_{j}\right)}_{\nabla_{\theta} \widehat{R_{u}}\left(f_{\theta}\right)}+\lambda \nabla_{\theta} \Omega(\theta) \\
\nabla_{\theta} h_{\theta}\left(u_{i}, u_{j}\right) & =-\left[\left(y_{i}-f_{\theta}\left(x_{i}\right)\right) \nabla_{\theta} f_{\theta}\left(x_{j}\right)+\left(y_{j}-f_{\theta}\left(x_{j}\right)\right) \nabla_{\theta} f_{\theta}\left(x_{i}\right)\right] k\left(z_{i}, z_{j}\right) .
\end{aligned}
$$

Asymptotic normality of $\nabla_{\theta} \widehat{R}_{U, \lambda}\left(f_{\theta^{*}}\right)$. We first show the asymptotic normality of $\nabla_{\theta} \widehat{R}_{U, \lambda}\left(f_{\theta^{*}}\right)$. We assume that there exists $z \in \mathcal{Z}$ such that $\mathbb{E}_{X}\left[\nabla_{\theta} f_{\theta^{*}}(X) \mid z\right] p(z) \neq 0$ or $\mathbb{E}_{X Y}\left[Y-f_{\theta^{*}}(X) \mid z\right] p(z) \neq 0$. Both terms being equal to zeros for all $z \in \mathcal{Z}$ leads to a singular $\nabla_{\theta}^{2} \widehat{R}_{U}\left(f_{\theta^{*}}\right)$ and the asymptotic distribution therefore becomes much more complicated to analyze.

Lemma 4. Suppose that $f_{\theta}$ and $\Omega(\theta)$ are first continuously differentiable about $\theta, \mathbb{E}\left[\left\|\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right\|_{2}^{2}\right]<$ $\infty$, there exists $z \in \mathcal{Z}$ such that $\mathbb{E}_{X}\left[\nabla_{\theta} f_{\theta^{*}}(X) \mid z\right] p(z) \neq 0$ or $\mathbb{E}_{X Y}\left[Y-f_{\theta^{*}}(X) \mid z\right] p(z) \neq 0$, and $\sqrt{n} \lambda \xrightarrow{\mathrm{P}} 0$. Then,

$$
\sqrt{n} \nabla_{\theta} \widehat{R}_{U, \lambda}\left(f_{\theta^{*}}\right) \xrightarrow{\mathrm{p}} N\left(\mathbf{0}, 4 \operatorname{diag}\left(\mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right)\right) .
$$

Proof. The proof follows from Serfling [1980, Section 5.5 .1 and Section 5.5.2] and we need to show that (i) $\nabla_{\theta} \widehat{R}_{U}\left(f_{\theta^{*}}\right) \xrightarrow{\mathrm{p}} \mathbf{0}$ and (ii) whether $\operatorname{Var}_{U}\left[\mathbb{E}_{U^{\prime}}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]>0$ or not. (i) can be obtained by the law of large numbers because $\nabla_{\theta} \widehat{R}_{U}\left(f_{\theta^{*}}\right)$ is a sample average of $\nabla_{\theta} R_{k}\left(f_{\theta^{*}}\right)=\mathbf{0}$.

To prove (ii), we first note that $\operatorname{Var}_{U}\left[\mathbb{E}_{U^{\prime}}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]=\mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]-$ $\underbrace{\mathbb{E}_{U u^{\prime}}^{2}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]}_{=0} \geq \mathbf{0}$, where equality holds if for any $U$, there is $\mathbb{E}_{U^{\prime}}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]=$ 0 , i.e.,

$$
\begin{aligned}
& \mathbb{E}_{U^{\prime}}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right] \\
& =-\mathbb{E}_{X^{\prime} Z^{\prime}}\left[\nabla_{\theta} f_{\theta^{*}}\left(X^{\prime}\right) k\left(Z^{\prime}, Z\right)\right]\left(Y-f_{\theta^{*}}(X)\right)-\mathbb{E}_{X^{\prime} Y^{\prime} Z^{\prime}}\left[\left(Y^{\prime}-f_{\theta^{*}}\left(X^{\prime}\right)\right) k\left(Z^{\prime}, Z\right)\right] \nabla_{\theta} f_{\theta^{*}}(X) \\
& =\mathbf{0} .
\end{aligned}
$$

As the above equation holds for any $Y$, the coefficient of $Y$ must be 0 :

$$
\mathbb{E}_{X^{\prime} Z^{\prime}}\left[\nabla_{\theta} f_{\theta^{*}}\left(X^{\prime}\right) k\left(Z^{\prime}, Z\right)\right]=\mathbb{E}_{Z^{\prime}}\left[\mathbb{E}_{X^{\prime}}\left[\nabla_{\theta} f_{\theta^{*}}\left(X^{\prime}\right) \mid Z^{\prime}\right] k\left(Z^{\prime}, Z\right)\right]=0,
$$

where we note that $\mathbb{E}\left[\nabla_{\theta} f_{\theta^{*}}\left(X^{\prime}\right) \mid Z^{\prime}\right] p\left(Z^{\prime}\right)=0$ for any $Z^{\prime}$ implied by the second function above. Similarly, the coefficient of $\nabla_{\theta} f_{\theta^{*}}(X)$ must be zero, which implies that $\mathbb{E}_{X^{\prime} Y^{\prime}}\left[\left(Y^{\prime}-f_{\theta^{*}}\left(X^{\prime}\right)\right) \mid Z^{\prime}\right] p\left(Z^{\prime}\right)=0$ for any $Z^{\prime}$. The two coefficients cannot be zero at the same time (otherwise against the given conditions), so $\operatorname{Var}_{U}\left[\mathbb{E}_{U^{\prime}}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]>$ 0 . Further due to the given condition $\mathbb{E}\left[\left\|\nabla_{\theta^{*}} h\left(U, U^{\prime}\right)\right\|_{2}^{2}\right]<\infty$, we obtain $\sqrt{n} \nabla_{\theta} \widehat{R}_{U}\left(f_{\theta^{*}}\right) \xrightarrow{\mathrm{p}}$ $N\left(\mathbf{0}, 4 \mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right)$ as per Serfling [1980, Section 5.5.1]. Finally, as $\sqrt{n} \lambda \xrightarrow{\mathrm{p}} 0$
and $\nabla_{\theta} \Omega\left(\theta^{*}\right)<\infty$ by the condition that $\Omega(\theta)$ is first continuously differentiable, we assert the conclusion by Slutsky's theorem,
$\sqrt{n} \nabla_{\theta} \widehat{R}_{U, \lambda}\left(f_{\theta^{*}}\right)=\sqrt{n} \nabla_{\theta} \widehat{R}_{U}\left(f_{\theta^{*}}\right)+\sqrt{n} \lambda \nabla_{\theta} \Omega\left(\theta^{*}\right) \xrightarrow{\mathrm{p}} N\left(\mathbf{0}, 4 \operatorname{diag}\left(\mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right)\right)$.
This concludes the proof.

Uniform consistency of $\nabla_{\theta}^{2} \widehat{R}_{U, \lambda}\left(f_{\theta}\right)$. Next, we consider the second derivative $\nabla_{\theta}^{2} \widehat{R}_{U, \lambda}\left(f_{\theta}\right)$ and show its uniform consistency. In what follows, we denote by $\|\cdot\|_{F}$ the Frobenius norm. We can express $\nabla_{\theta}^{2} \widehat{R}_{U, \lambda}\left(f_{\theta}\right)$ as

$$
\begin{aligned}
\nabla_{\theta}^{2} \widehat{R}_{U, \lambda}\left(f_{\theta}\right)= & \underbrace{\frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j \neq i}^{n} \nabla_{\theta}^{2} h_{\theta}\left(u_{i}, u_{j}\right)}_{\nabla_{\theta} \hat{R} \widehat{L}\left(f_{\theta}\right)}+\lambda \nabla_{\theta}^{2} \Omega(\theta) \\
\nabla_{\theta}^{2} h_{\theta}\left(u_{i}, u_{j}\right)= & {\left[\nabla_{\theta} f_{\theta}\left(x_{i}\right) \nabla_{\theta} f_{\theta}^{\top}\left(x_{j}\right)-\left(y_{i}-f_{\theta}\left(x_{i}\right)\right) \nabla_{\theta}^{2} f_{\theta}\left(x_{j}\right)\right.} \\
& \left.\quad+\nabla_{\theta} f_{\theta}\left(x_{j}\right) \nabla_{\theta} f_{\theta}^{\top}\left(x_{i}\right)-\left(y_{j}-f_{\theta}\left(x_{j}\right)\right) \nabla_{\theta}^{2} f_{\theta}\left(x_{i}\right)\right] k\left(z_{i}, z_{j}\right) .
\end{aligned}
$$

Lemma 5. Suppose that $f_{\theta}$ and $\Omega(\theta)$ are twice continuously differentiable about $\theta, \Theta$ is compact, $\mathbb{E}\left[\left|f_{\theta}(X)\right|\right]<\infty, \mathbb{E}\left[\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}\right]<\infty, \mathbb{E}\left[\left\|\nabla_{\theta}^{2} f_{\theta}(X)\right\|_{F}\right]<\infty, \mathbb{E}[|Y|]<\infty$, $\lambda \xrightarrow{\mathrm{p}} 0$ and Assumption 1 holds. Then, $\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]$ is continuous about $\theta$ and

$$
\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} \widehat{R}_{U, \lambda}\left(f_{\theta}\right)-\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]\right\|_{\mathrm{F}} \xrightarrow{\mathrm{p}} 0 .
$$

Proof. The proof is similar to that of Lemma 3 and both applies extended results of Newey and McFadden [1994, Lemma 2.4]. As $\left(u_{i}, u_{j}\right)_{i \neq j}$ being strictly stationary in the sense of Newey and McFadden [1994, Footnote 18 in P.2129] has been shown in Lemma 3, we only need to show that (i) $\nabla_{\theta}^{2} h_{\theta}\left(u, u^{\prime}\right)$ is continuous at each $\theta \in \Theta$ with probability one and (ii) there exists $d\left(u, u^{\prime}\right) \geq\left\|\nabla_{\theta}^{2} h_{\theta}\left(u, u^{\prime}\right)\right\|_{F}$ for all $\theta \in \Theta$ and $\mathbb{E}\left[d\left(U, U^{\prime}\right)\right]<\infty$. We exploit the triangle inequality of the Frobenius norm and obtain

$$
\begin{aligned}
& \left\|\nabla_{\theta}^{2} h_{\theta}\left(u, u^{\prime}\right)\right\|_{\mathrm{F}} \\
& \leq\left[2\left\|\nabla_{\theta} f_{\theta}(x) \nabla_{\theta} f_{\theta}^{\top}\left(x^{\prime}\right)\right\|_{F}+\left(|y|+\left|f_{\theta}(x)\right|\right)\left\|\nabla_{\theta}^{2} f_{\theta}\left(x^{\prime}\right)\right\|_{\mathrm{F}}+\left(\left|y^{\prime}\right|+\left|f_{\theta}\left(x^{\prime}\right)\right|\right)\left\|\nabla_{\theta}^{2} f_{\theta}(x)\right\|_{\mathrm{F}}\right] k\left(z, z^{\prime}\right) \\
& \quad \equiv d\left(u, u^{\prime}\right)
\end{aligned}
$$

We first show $d\left(u, u^{\prime}\right)$ is bounded for bounded $u, u^{\prime}$. As $f_{\theta}$ is twice continuously differentiable about $\theta$ and $\Theta$ is compact, we have $f_{\theta}(x)$ bounded as well as each entry of $\nabla_{\theta} f_{\theta}(x)$ and $\nabla_{\theta}^{2} f_{\theta}(x)$ for $\|x\|<\infty$. Further taking into account that $k(\cdot, \cdot)$ is bounded as per Assumption 1, we know that $d\left(u, u^{\prime}\right)<\infty$ if $u, u^{\prime}$ are bounded, and it follows that (i) $\nabla_{\theta}^{2} h_{\theta}\left(u, u^{\prime}\right)$ is continuous at each $\theta \in \Theta$ with probability one as $f_{\theta}$ is twice continuously differentiable.

We then show that (ii) $\mathbb{E}_{U, U^{\prime}}\left[d\left(U, U^{\prime}\right)\right]<\infty$ by the following inequalities:

$$
\begin{aligned}
& \mathbb{E}_{U, U^{\prime}}\left[d\left(U, U^{\prime}\right)\right] \\
& \leq 2 \mathbb{E}\left[\left\|\nabla_{\theta} f_{\theta}(X) \nabla_{\theta} f_{\theta}^{\top}\left(X^{\prime}\right)\right\|_{F}+\left(|Y|+\left|f_{\theta}(X)\right|\right)\left\|\nabla_{\theta}^{2} f_{\theta}\left(X^{\prime}\right)\right\|_{F}\right] \sup _{z} k(z, z) \\
& =2(\underbrace{\mathbb{E}\left[\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{F}\right]}_{<\infty} \mathbb{E}\left[\left\|\nabla_{\theta} f_{\theta}^{\top}\left(X^{\prime}\right)\right\|_{F}\right]+\underbrace{\mathbb{E}\left[|Y|+\left|f_{\theta}(X)\right|\right]}_{<\infty} \underbrace{\mathbb{E}\left[\left\|\nabla_{\theta}^{2} f_{\theta}\left(X^{\prime}\right)\right\|_{F}\right]}_{<\infty}) \sup _{z} k(z, z)
\end{aligned}
$$

$<\infty$.
Therefore, we obtain $\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} \widehat{R}_{U}\left(f_{\theta}\right)-\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]\right\|_{\mathrm{F}} \xrightarrow{\mathrm{p}} 0$ following from the extended results in the remarks of Newey and McFadden [1994, Lemma 2.4]. Furthermore, from the conditions that $\Omega(\theta)$ is twice continuously differentiable and the parameter space $\Theta$ is compact, we obtain that $\left\|\nabla_{\theta}^{2} \Omega(\theta)\right\|_{\mathrm{F}}<\infty$ for any $\theta \in \Theta$. Finally, it follows from the Slutsky's theorem that

$$
\begin{aligned}
& \sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} \widehat{R}_{U, \lambda}\left(f_{\theta}\right)-\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]\right\|_{\mathrm{F}} \\
& \leq \sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} \widehat{R}_{U}\left(f_{\theta}\right)-\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]\right\|_{\mathrm{F}}+\lambda \sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} \Omega(\theta)\right\|_{\mathrm{F}} \xrightarrow{\mathrm{p}} 0 .
\end{aligned}
$$

This concludes the proof.
Theorem 16 (Asymptotic normality of $\left.\hat{\theta}_{U}\right)$. Suppose that $H=\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]$ is nonsingular, $\Theta$ compact, $\mathbb{E}\left[\left|f_{\theta}(X)\right|\right]<\infty, \mathbb{E}[|Y|]<\infty, f_{\theta}$ and $\Omega(\theta)$ are twice continuously differentiable about $\theta, \mathbb{E}\left[\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}\right]<\infty, \mathbb{E}\left[\left\|\nabla_{\theta}^{2} f_{\theta}(X)\right\|_{F}\right]<\infty, \sqrt{n} \lambda \xrightarrow{\mathrm{P}} 0, R_{k}\left(f_{\theta}\right)$ is uniquely minimized at $\theta^{*}$ which is an interior point of $\Theta, \mathbb{E}\left[\left\|\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right\|_{2}^{2}\right]<\infty$ and Assumptions 1 hold. Then

$$
\sqrt{n}\left(\hat{\theta}_{U}-\theta^{*}\right) \xrightarrow{\mathrm{p}} N\left(\mathbf{0}, 4 H^{-1} \operatorname{diag}\left(\mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right) H^{-1}\right) .
$$

Proof. The proof follows by Newey and McFadden [1994, Theorem 3.1] and we need to show that (i) $\widehat{\theta}_{U} \xrightarrow{\mathrm{p}} \theta^{*}$; (ii) $\widehat{R}_{U, \lambda}(\theta)$ is twice continuously differentiable; (iii) $\sqrt{n} \nabla_{\theta} \widehat{R}_{U, \lambda}\left(f_{\theta^{*}}\right) \xrightarrow{\mathrm{p}} N\left(\mathbf{0}, 4 \mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right) ;$ (iv) there is $H(\theta)$ that is continuous at $\theta^{*}$ and $\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} \widehat{R}_{U, \lambda}\left(f_{\theta}\right)-H(\theta)\right\|_{F} \xrightarrow{\mathrm{p}} 0$; (v) $H\left(\theta^{*}\right)$ is non-singular.

The proof of (i) is very similar to Theorem 15 except that we consider finite dimensional parameter space instead of functional space. For a neat proof, we would like to omit the detailed proof here. We can first show the uniform consistency $\sup _{\theta \in \Theta}\left|\widehat{R}_{U, \lambda}\left(f_{\theta}\right)-R_{k}\left(f_{\theta}\right)\right| \xrightarrow{\mathrm{p}} 0$ and $R_{k}\left(f_{\theta}\right)$ is continuous about $\theta$ similarly to Lemma 3. Here, the proof is based on the conditions $\mathbb{E}[|Y|]<\infty, \Theta$ is compact, $\mathbb{E}\left[\left|f_{\theta}(X)\right|\right]<\infty$ and $f_{\theta}$ is twice continuously differentiable about $\theta$, and Assumption 1 holds. Then, $\hat{\theta}_{U} \xrightarrow{\mathrm{p}} \theta^{*}$ similarly to Theorem 15 , because of the extra condition $R_{k}\left(f_{\theta}\right)$ is uniquely minimized at $\theta^{*}$.

Furthermore, from the conditions that $\Theta$ is compact, $f_{\theta}$ is twice continuously differentiable about $\theta, \mathbb{E}\left[\left|f_{\theta}(X)\right|\right]<\infty, \mathbb{E}\left[\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}\right]<\infty, \mathbb{E}\left[\left\|\nabla_{\theta}^{2} f_{\theta}(X)\right\|_{F}\right]<\infty$, $\mathbb{E}[|Y|]<\infty$ and $k\left(z, z^{\prime}\right)$ is bounded as implied by Assumption 1, we can obtain (ii) $\widehat{R}_{V, \lambda}(\theta)$ is twice continuously differentiable about $\theta$. Given $H=\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]=$ $\nabla_{\theta}^{2} R_{k}(\theta)$ is non-singular and $R_{k}\left(f_{\theta}\right)$ is uniquely minimized at $\theta^{*}$, we can obtain that the Hessian matrix $H$ is positive definite,

$$
H=2 \mathbb{E}_{(X Y Z),\left(X^{\prime} Y^{\prime} Z^{\prime}\right)}\left[\left(\nabla_{\theta} f_{\theta^{*}}(X) \nabla_{\theta} f_{\theta^{*}}^{\top}\left(X^{\prime}\right)-\left(Y-f_{\theta^{*}}(X)\right) \nabla_{\theta}^{2} f_{\theta^{*}}\left(X^{\prime}\right)\right) k\left(Z, Z^{\prime}\right)\right] \succ \mathbf{0} .
$$

If for all $z \in \mathcal{Z}$, there is $\mathbb{E}_{X}\left[\nabla_{\theta} f_{\theta^{*}}(X) \mid z\right] p(z)=\mathbf{0}$ and $\mathbb{E}_{X Y}\left[Y-f_{\theta^{*}}(X) \mid z\right] p(z)=0$, then we can see that the above function $H=\mathbf{0}$ which contradicts $H \succ \mathbf{0}$. Therefore, there must exist $z$ s.t. $\mathbb{E}_{X}\left[\nabla_{\theta} f_{\theta^{*}}(X) \mid z\right] p(z) \neq \mathbf{0}$ or $\mathbb{E}_{X Y}\left[Y-f_{\theta^{*}}(X) \mid z\right] p(z) \neq 0$. Then, it follows by Lemma 4 that (iii) $\sqrt{n} \nabla_{\theta} \widehat{R}_{U, \lambda}\left(f_{\theta^{*}}\right) \xrightarrow{\mathrm{p}} N\left(\mathbf{0}, 4 \mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right)$.

Finally by Lemma 5 , we know that $H(\theta)=\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]$ and $H\left(\theta^{*}\right)=H$, so (iv) and (v) are satisfied. Now, conditions of Newey and McFadden [1994, Theorem 3.1] are all satisfied, so we assert the conclusion.

## D.2.10 Proof of Theorem 9

We restate the notations

$$
\begin{aligned}
& \widehat{R}_{V, \lambda}\left(f_{\theta}\right):=\underbrace{\frac{1}{n^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n} h_{\theta}\left(u_{i}, u_{j}\right)}_{\widehat{R}_{V}\left(f_{\theta}\right)}+\lambda \Omega(\theta) \\
& h_{\theta}\left(u_{i}, u_{j}\right):=\left(y_{i}-f_{\theta}\left(x_{i}\right)\right) k\left(z_{i}, z_{j}\right)\left(y_{j}-f_{\theta}\left(x_{j}\right)\right),
\end{aligned}
$$

Lemma 6. Suppose that conditions of Lemma 4 hold. Then

$$
\sqrt{n} \nabla_{\theta} \widehat{R}_{V, \lambda}\left(f_{\theta^{*}}\right) \xrightarrow{\mathrm{p}} N\left(\mathbf{0}, 4 \mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right) .
$$

Proof. As $\mathbb{E}\left[\left\|\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right\|_{2}^{2}\right]<\infty, \sqrt{n} \nabla_{\theta} \widehat{R}_{V}\left(f_{\theta^{*}}\right)$ has the same limit distribution as that of $\sqrt{n} \nabla_{\theta} \widehat{R}_{U}\left(f_{\theta^{*}}\right)$ by Serfling [1980, Section 5.7.3]. Furthermore, by $\sqrt{n} \lambda \xrightarrow{p} 0$ and $\nabla_{\theta} \Omega\left(\theta^{*}\right)<\infty$ from that $\Omega(\theta)$ is first continuously differentiable, we assert the conclusion by Slutsky's theorem

$$
\sqrt{n} \nabla_{\theta} \widehat{R}_{V, \lambda}\left(f_{\theta^{*}}\right)=\sqrt{n} \nabla_{\theta} \widehat{R}_{V}+\sqrt{n} \lambda \nabla_{\theta} \Omega\left(\theta^{*}\right) \xrightarrow{\mathrm{p}} N\left(\mathbf{0}, 4 \mathbb{E}_{U}\left[\mathbb{E}_{U^{\prime}}^{2}\left[\nabla_{\theta} h_{\theta^{*}}\left(U, U^{\prime}\right)\right]\right]\right) .
$$

Lemma 7. Suppose that $f_{\theta}$ and $\Omega(\theta)$ are twice continuously differentiable about $\theta, \Theta$ is compact, $\mathbb{E}\left[\sup _{\theta \in \Theta}\left|f_{\theta}(X)\right|^{2}\right]<\infty, \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}^{2}\right]<\infty, \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} f_{\theta}(X)\right\|_{F}^{2}\right]<$ $\infty, \mathbb{E}\left[|Y|^{2}\right]<\infty, \lambda \xrightarrow{\mathrm{p}} 0$ and Assumption 1 holds. Then, $\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]$ is continuous about $\theta$ and $\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} \widehat{R}_{V}\left(f_{\theta}\right)-\mathbb{E}\left[\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right]\right\|_{\mathrm{F}} \xrightarrow{\mathrm{p}} 0$.

Proof. We apply Newey and McFadden [1994, Lemma 8.5] for this proof and need to show (i) $\nabla_{\theta}^{2} h_{\theta}\left(u, u^{\prime}\right)$ is continuous about each $\theta \in \Theta$ with probability one, and (ii) $\mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right\|_{\mathrm{F}}\right]<\infty$ and $\mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} h_{\theta}(U, U)\right\|_{\mathrm{F}}\right]<\infty$.

We first see that $\left\|\nabla_{\theta}^{2} h_{\theta}\left(u, u^{\prime}\right)\right\|_{\mathrm{F}}$ and $\left\|\nabla_{\theta}^{2} h_{\theta}(u, u)\right\|_{\mathrm{F}}$ are bounded for finite $u, u^{\prime}$ because $f_{\theta}$ is twice continuously differentiable about $\theta$. It follows that (i) $\nabla_{\theta}^{2} h_{\theta}\left(u, u^{\prime}\right)$ is continuous about $\theta$ with probability one. We then derive upper bounds for $\mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right\|_{\mathrm{F}}\right]$ and $\mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} h_{\theta}(U, U)\right\|_{\mathrm{F}}\right]$ so as to show their boundedness,

$$
\begin{aligned}
& \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} h_{\theta}\left(U, U^{\prime}\right)\right\|_{\mathrm{F}}\right] \\
& \leq 2 \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X) \nabla_{\theta} f_{\theta}^{\top}\left(X^{\prime}\right)\right\|_{F}+\left(|Y|+\left|f_{\theta}(X)\right|\right)\left\|\nabla_{\theta}^{2} f_{\theta}\left(X^{\prime}\right)\right\|_{\mathrm{F}}\right] \sup _{z} k(z, z) \\
& \leq 2 \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}\right]^{2}+2 \mathbb{E}\left[|Y|+\sup _{\theta \in \Theta}\left|f_{\theta}(X)\right|\right] \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} f_{\theta}\left(X^{\prime}\right)\right\|_{\mathrm{F}}\right] \sup _{z} k(z, z) \\
& <\infty,
\end{aligned}
$$

and

$$
\begin{aligned}
& \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} h_{\theta}(U, U)\right\|_{\mathrm{F}}\right] \\
& \leq 2 \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X) \nabla_{\theta} f_{\theta}^{\top}(X)\right\|_{F}+\left(|Y|+\left|f_{\theta}(X)\right|\right)\left\|\nabla_{\theta}^{2} f_{\theta}(X)\right\|_{\mathrm{F}}\right] \sup _{z} k(z, z) \\
& \leq 2 \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}^{2}\right]+2 \mathbb{E}\left[\left(|Y|+\sup _{\theta \in \Theta}\left|f_{\theta}(X)\right|\right)^{2}\right] \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} f_{\theta}\left(X^{\prime}\right)\right\|_{\mathrm{F}}^{2}\right] \sup _{z} k(z, z) \\
& \leq 2 \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}^{2}\right]+\mathbb{E}\left[\left(|Y|+\sup _{\theta \in \Theta}\left|f_{\theta}(X)\right|\right)^{2}\right]+\mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} f_{\theta}\left(X^{\prime}\right)\right\|_{\mathrm{F}}^{2}\right] \sup _{z} k(z, z) \\
& \leq 2 \mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta} f_{\theta}(X)\right\|_{2}^{2}\right]+2 \mathbb{E}\left[|Y|^{2}+\sup _{\theta \in \Theta}\left|f_{\theta}(X)\right|^{2}\right]+\mathbb{E}\left[\sup _{\theta \in \Theta}\left\|\nabla_{\theta}^{2} f_{\theta}\left(X^{\prime}\right)\right\|_{\mathrm{F}}^{2}\right] \sup _{z} k(z, z)
\end{aligned}
$$

$<\infty$.

Thus, we assert the conclusion by Newey and McFadden [1994, Lemma 8.5].

Proof of Theorem 9. The proof is the same as that of Theorem 16 except that $\widehat{R}_{U}$ is replaced by $\widehat{R}_{V}$.

## D.2.11 Asymptotic Normality in the Infinite-dimension Case

We firstly state the asymptotic normality theorem for $\hat{f}_{U}$ and its proof. Afterwards, we provide the proof of Theorem 10 whose proof is a slightly modified version of that of $\hat{f}_{u}$.

Theorem 17. Suppose Assumption 1 holds, $l$ is a bounded kernel, $k$ is a uniformly bounded function, and $\lambda \geq \lambda_{0}$ holds. Also, suppose that $\mathcal{X}, \mathcal{Z}$, and $\mathcal{Y}$ are compact spaces, and there exists $s \in(0,2)$ and a constant $C_{H}>0$ such that $\log \mathcal{N}\left(\varepsilon, \mathcal{H}_{l},\|\cdot\|_{L^{\infty}}\right) \leq C_{H} \varepsilon^{-s}$ for any $\varepsilon \in(0,1)$. If $\lambda-\lambda_{0}=o\left(n^{-1 / 2}\right)$ holds, then there exists a Gaussian process $\mathbb{G}_{P}^{*}$ such that

$$
\sqrt{n}\left(\hat{f}_{U}-f_{\lambda_{0}}^{*}\right) \rightsquigarrow \mathbb{G}_{P}^{*} \text { in } \mathcal{H}_{l} .
$$

An exact covariance of $G_{P}^{*}$ is described in the proof. The proof is based on the uniform convergence of U-processes on the function space [Arcones and Gine, 1993] and the functional delta method using the asymptotic expansion of the loss function [Hable, 2012]. This asymptotic normality allows us to perform statistical inference, such as tests, even in the non-parametric case.

We discuss the boundedness and covering number assumptions in Theorem 10 and Theorem 17. For the boundedness assumption, many common kernels, such as the Gaussian RBF kernel, the Laplacian kernel, and the Mercer kernel, satisfy it. For the covering number, the common kernels above also satisfy it with a certain parameter configuration. For example, the Gaussian kernel (see Section 4 in Steinwart and Christmann [2008]) and the Mercer kernel (explained in Zhou [2002]).

To prove the theorem, we provide some notation. Let $P$ be a probability measure which generates $u=(x, y, z)$ and $\mathcal{W}=\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$. Also, we define a function $h_{f}\left(u, u^{\prime}\right)=(y-f(x))\left(y^{\prime}-f\left(x^{\prime}\right)\right) k\left(z, z^{\prime}\right)$. Let $\mathbb{H}:=\left\{h_{f}: \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R} \mid f \in \mathcal{H}_{l}\right\}$. For preparation, we define $P^{1} h_{f}: \mathcal{W} \rightarrow \mathbb{R}$ as $P^{1} h_{f}(\cdot)=\left(\int h_{f}(u, \cdot)+h_{f}(\cdot, u) d P(u)\right) / 2$ for $h_{f} \in \mathbb{H}$. For a signed measure $Q$ on $\mathcal{W}$, we define a measure $Q^{2}:=Q \otimes Q$ on $\mathcal{W} \times \mathcal{W}$. Then, we can rewrite the U-statistic risk as

$$
\widehat{R}_{U}(f)=\frac{(n-2)!}{n!} \sum_{i=1}^{n} \sum_{j \neq i}^{n} h_{f}\left(u_{i}, u_{j}\right)=: U_{n}^{2} h_{f},
$$

where $U_{n}$ is an empirical measure for the $U$-statistics. Similarly, we can rewrite the V-statistic risk as

$$
\widehat{R}_{V}(f)=\frac{1}{n^{2}} \sum_{i=1}^{n} \sum_{j=1}^{n} h_{f}\left(u_{i}, u_{j}\right)=: P_{n}^{2} h_{f}
$$

where $P_{n}$ is an empirical measure of $u$.
Further, we define a functional associated with measure. We consider functional spaces $\mathcal{G}_{1}:=\{g: \mathcal{W} \times \mathcal{W} \rightarrow[0,1] \mid$ a convex set $\omega$ s.t. $g=\mathbf{1}\{\cdot \leq \omega\}\}$ and $\mathcal{G}_{2}:=$ $\left\{g: \mathcal{W} \times \mathcal{W} \rightarrow[0,1] \mid \exists f, f^{\prime} \in \mathcal{H}_{l,}, g\left(u, u^{\prime}\right)=h_{f}\left(u, u^{\prime}\right)\left(f^{\prime}(x)+f\left(x^{\prime}\right)\right)\right\}$. Note that $\mathcal{G}_{1}$ contains a functional which corresponds to the $U_{n}^{2}$. Then, we consider a set of
functionals
$B_{S}:=\left\{F: \mathcal{G}_{1} \cup \mathcal{G}_{2} \rightarrow \mathbb{R} \mid \exists\right.$ non-zero finite $Q^{2}$ s.t. $\left.F(g)=\int g d Q^{2}, \quad \forall g \in \mathcal{G}_{1} \cup \mathcal{G}_{2}\right\}$, and also let $B_{0}$ be the closed linear span of $B_{S}$. For a functional $F \in B_{S}$, let $\iota(F)$ be a measure which satisfies

$$
F(g)=\int g d \iota(F)
$$

Uniform Central Limit Theorem: We firstly achieve the uniform convergence. Note that a measure $P^{2}$ satisfies $P^{2} h_{f}:=\mathbb{E}_{\left(U, U^{\prime}\right)}\left[h_{f}\left(U, U^{\prime}\right)\right]$. The convergence theorem for U-processes is as follows:
Theorem 18 (Theorem 4.4 in Arcones and Gine [1993]). Suppose $\mathbb{H}$ is a set of uniformly bounded class and symmetric functions, such that $\left\{P^{1} h_{f} \mid h_{f} \in \mathbb{H}\right\}$ is a Donsker class and

$$
\lim _{n \rightarrow \infty} \mathbb{E}\left[n^{-1 / 2} \log \mathcal{N}\left(n^{-1 / 2} \varepsilon, \mathbb{H},\|\cdot\|_{L^{1}\left(U_{n}^{2}\right)}\right)\right]=0
$$

holds, for all $\varepsilon>0$. Then, we obtain

$$
\sqrt{n}\left(U_{n}^{2}-P^{2}\right) \rightsquigarrow 2 \mathbb{G}_{p^{1}}, \text { in } \ell^{\infty}(\mathbb{H}) .
$$

Here, $\mathbb{G}_{P^{1}}$ denotes a Brownian bridge, which is a Gaussian process on $\mathbb{H}$ with zero mean and a covariance

$$
\mathbb{E}_{U}\left[P^{1} h_{f}(U) P^{1} h_{f}^{\prime}(U)\right]-\mathbb{E}_{U}\left[P^{1} h_{f}(U)\right] \mathbb{E}_{U}\left[P^{1} h_{f}^{\prime}(U)\right],
$$

with $h_{f}, h_{f}^{\prime} \in \mathbb{H}$.
To apply the theorem, we have to show that $\mathbb{H}$ satisfies the condition in Theorem 18. We firstly provide the following bound:
Lemma 8. For any $f, f^{\prime} \in \mathcal{H}_{l}$ such that $\|f\|_{L^{\infty}} \vee\left\|f^{\prime}\right\|_{L^{\infty}} \leq B$ holds, $y, y^{\prime} \in[-B, B]$ and $u, u^{\prime} \in \mathcal{W}$, we have

$$
\left|h_{f}\left(u, u^{\prime}\right)-h_{f^{\prime}}\left(u, u^{\prime}\right)\right| \leq 4 B\left|k\left(z, z^{\prime}\right)\right|\left\|f-f^{\prime}\right\|_{L^{\infty}} .
$$

Proof of Lemma 8. We simply obtain the following:

$$
\begin{aligned}
& \left|h_{f}\left(u, u^{\prime}\right)-h_{f^{\prime}}\left(u, u^{\prime}\right)\right| \\
& =\left|(y-f(x)) k\left(z, z^{\prime}\right)\left(y^{\prime}-f\left(x^{\prime}\right)\right)-\left(y-f^{\prime}(x)\right) k\left(z, z^{\prime}\right)\left(y^{\prime}-f^{\prime}\left(x^{\prime}\right)\right)\right| \\
& =\left|k\left(z, z^{\prime}\right)\right|\left|(y-f(x))\left(y^{\prime}-f\left(x^{\prime}\right)\right)-\left(y-f^{\prime}(x)\right)\left(y^{\prime}-f^{\prime}\left(x^{\prime}\right)\right)\right| \\
& =\left|k\left(z, z^{\prime}\right)\right|\left|y^{\prime}\left(f^{\prime}(x)-f(x)\right)+y\left(f^{\prime}\left(x^{\prime}\right)-f\left(x^{\prime}\right)\right)+f(x) f\left(x^{\prime}\right)-f^{\prime}(x) f^{\prime}\left(x^{\prime}\right)\right| \\
& =\left|k\left(z, z^{\prime}\right)\right|\left|y^{\prime}\left(f^{\prime}(x)-f(x)\right)+y\left(f^{\prime}\left(x^{\prime}\right)-f\left(x^{\prime}\right)\right)+f\left(x^{\prime}\right)\left(f(x)-f^{\prime}(x)\right)-f^{\prime}(x)\left(f\left(x^{\prime}\right)-f^{\prime}\left(x^{\prime}\right)\right)\right| \\
& \leq\left|k\left(z, z^{\prime}\right)\right|\left\{\left|y^{\prime}\right|\left|f^{\prime}(x)-f(x)\right|+|y|\left|f^{\prime}\left(x^{\prime}\right)-f\left(x^{\prime}\right)\right|+\left|f\left(x^{\prime}\right)\right|\left|f(x)-f^{\prime}(x)\right|-\left|f^{\prime}(x)\right|\left|f\left(x^{\prime}\right)-f^{\prime}\left(x^{\prime}\right)\right|\right\} \\
& \leq 4 B\left|k\left(z, z^{\prime}\right)\right|| | f-f^{\prime} \|_{L^{\infty},}
\end{aligned}
$$

as required.
Lemma 9. Suppose the assumptions of Theorem 10 hold. Then, the followings hold:

1. $\left\{P^{1} h_{f} \mid h_{f} \in \mathbb{H}\right\}$ is a Donsker class.
2. For any $\varepsilon>0$, the following holds:

$$
\lim _{n \rightarrow \infty} E\left[n^{-1 / 2} \log \mathcal{N}\left(n^{-1 / 2} \varepsilon, \mathbb{H},\|\cdot\|_{L^{1}\left(U_{n}^{2}\right)}\right)\right]=0 .
$$

Proof of Lemma 9. For preparation, fix $\varepsilon>0$ and set $N=\mathcal{N}\left(\varepsilon, \mathcal{H}_{l},\|\cdot\|_{L^{\infty}}\right)$. Also, let $Q$ be an arbitrary finite discrete measure. Then, by the definition of a bracketing number, there exist $N$ functions $\left\{f_{i} \in \mathcal{H}_{l}\right\}_{i=1}^{N}$ such that for any $f \in \mathcal{H}_{l}$ there exists $i \in\{1,2, \ldots, N\}$ such as $\left\|f-f_{i}\right\|_{L^{\infty}} \leq \varepsilon$.

For the first condition, as shown in Equation (2.1.7) in Van der Vaart and Wellner [1996], it is sufficient to show

$$
\sup _{Q} \log \mathcal{N}\left(\varepsilon,\left\{P^{1} h_{f} \mid h_{f} \in \mathbb{H}\right\},\|\cdot\|_{L^{2}(Q)}\right) \leq c \varepsilon^{-2 \delta}(\varepsilon \rightarrow 0)
$$

for arbitrary $\delta \in(0,1)$. Here, $c>0$ is some constant, and $Q$ is taken from all possible finite discrete measure. To this end, it is sufficient to show that $\log \mathcal{N}\left(\varepsilon,\left\{P^{1} h_{f} \mid h_{f} \in\right.\right.$ $\left.\mathbb{H}\},\|\cdot\|_{L^{2}(Q)}\right) \leq c^{\prime} \log \mathcal{N}\left(\varepsilon, \mathcal{H}_{l},\|\cdot\|_{L^{\infty}}\right)$ with a constant $c^{\prime}>0$. Fix $P^{1} h_{f} \in\left\{P^{1} h_{f} \mid\right.$ $\left.h_{f} \in \mathbb{H}\right\}$ arbitrary, and set $f_{i}$ which satisfies $\left\|f-f_{i}\right\|_{L^{2}(Q)} \leq \varepsilon$. Then, we have

$$
\begin{aligned}
& \left\|P^{1} h_{f}-P^{1} h_{f_{i}}\right\|_{L^{2}(Q)}^{2} \\
& =\int\left\{\int\left(h_{f}\left(u, u^{\prime}\right)+h_{f}\left(u^{\prime}, u\right)\right) / 2-\left(h_{f_{i}}\left(u, u^{\prime}\right)+h_{f_{i}}\left(u^{\prime}, u\right)\right) / 2 d P\left(u^{\prime}\right)\right\}^{2} d Q(u) \\
& =\int\left(\int h_{f}\left(u, u^{\prime}\right)-h_{f_{i}}\left(u, u^{\prime}\right) d P\left(u^{\prime}\right)\right)^{2} d Q(u) \\
& \leq C \int\left(\int\left|k\left(z, z^{\prime}\right)\right| d P\left(u^{\prime}\right)\right)^{2} d Q(u)\left\|f-f_{i}\right\|_{L^{\infty}}^{2} \\
& \leq C^{\prime}\left\|f-f_{i}\right\|_{L^{\infty}}^{2} \\
& \leq C^{\prime} \varepsilon^{2}
\end{aligned}
$$

with constants $C, C^{\prime}>0$. The first inequality follows Lemma 8 with the bounded property of $f, f^{\prime}$ and $\mathcal{Y}$. The second inequality follows the bounded condition of $k$ in Theorem 17. Hence, the entropy condition shows the first statement.

For the second condition, we have the similar strategy. For any $h_{f} \in \mathbb{H}$, we consider $i \in\{1,2, \ldots, N\}$ such that $\left\|f-f_{i}\right\|_{L^{\infty}} \leq \varepsilon$. Then, we measure the following value

$$
\left\|h_{f}-h_{f_{i}}\right\|_{L^{1}\left(u_{n}^{2}\right)}=\int\left|h_{f}\left(u, u^{\prime}\right)-h_{f_{i}}\left(u, u^{\prime}\right)\right| d U_{n}^{2}\left(u, u^{\prime}\right) \leq C^{\prime \prime}\left\|f-f_{i}\right\|_{L^{\infty}} \leq C^{\prime \prime} \varepsilon,
$$

with a constant $C^{\prime \prime}>0$. Hence, we have

$$
\begin{aligned}
\mathbb{E}\left[n^{-1 / 2} \log \mathcal{N}\left(n^{-1 / 2} \varepsilon, \mathbb{H},\|\cdot\|_{L^{1}\left(U_{n}^{2}\right)}\right)\right] & \leq n^{-1 / 2} \log \mathcal{N}\left(n^{-1 / 2} \varepsilon, \mathcal{H}_{l},\|\cdot\|_{L^{\infty}}\right) \\
& \leq C n^{-1 / 2}\left(\frac{n^{1 / 2}}{\varepsilon}\right)^{s}=C n^{(s-1) / 2} \rightarrow 0,(n \rightarrow \infty)
\end{aligned}
$$

since $s \in(0,1)$.
From Theorem 18 and Lemma 9, we rewrite the central limit theorem utilizing terms of functionals. Note that $\iota^{-1}\left(U_{n}^{2}\right), \iota^{-1}\left(P^{2}\right) \in B_{S}$ holds. Then, we can obtain

$$
\sqrt{n}\left(\iota^{-1}\left(U_{n}^{2}\right)-\iota^{-1}\left(P^{2}\right)\right) \rightsquigarrow 2 \mathbb{G}_{P^{1}} \text { in } \ell^{\infty}(\mathbb{H}) .
$$

Learning Map and Functional Delta Method: We consider a learning map $S$ : $B_{S} \rightarrow \mathcal{H}_{l}$. For a functional $F \in B_{S}$, we define

$$
S_{\lambda}(F):=\underset{f \in \mathcal{H}_{l}}{\operatorname{argmin}} \iota(F) h_{f}+\lambda\|f\|_{\mathcal{H}_{l}}^{2} .
$$

Obviously, we have

$$
\hat{f}=S_{\lambda}\left(\iota^{-1}\left(U_{n}^{2}\right)\right), \text { and } f_{\lambda_{0}}^{*}=S_{\lambda_{0}}\left(\iota^{-1}\left(P^{2}\right)\right) .
$$

We consider a derivative of $S_{\lambda}$ in the sense of the Gateau differentiation by the following steps.

Firstly, we define a partial derivative of the map $R_{Q^{2}}(f)$. To investigate the optimality of the minimizer of

$$
R_{Q^{2}, \lambda}(f):=\int h_{f}\left(u, u^{\prime}\right) d Q^{2}\left(u, u^{\prime}\right)+\lambda\|f\|_{\mathcal{H}_{l}}^{2} .
$$

To this end, we consider the following derivative $\nabla R_{Q^{2}, \lambda}[f]: \mathcal{H}_{l} \rightarrow \mathcal{H}_{l}$ with a direction $f$ as

$$
\nabla R_{Q^{2}, \lambda}[f]\left(f^{\prime}\right):=2 \lambda f+\int \partial_{f, 1} h_{f}\left(u, u^{\prime}\right) f^{\prime}(x)+\partial_{f, 2} h_{f}\left(u, u^{\prime}\right) f^{\prime}\left(x^{\prime}\right) d Q^{2}\left(u, u^{\prime}\right) .
$$

Here, $\partial_{f, 1} h_{f}$ is a partial derivative of $h_{f}$ in terms of the input $f(x)$ as

$$
\partial_{f, 1} h_{f}\left(u, u^{\prime}\right)=\partial_{t \mid t=f(x)}(y-t) k\left(z, z^{\prime}\right)\left(y^{\prime}-f\left(x^{\prime}\right)\right)=-\left(y^{\prime}-f\left(x^{\prime}\right)\right) k\left(z, z^{\prime}\right)
$$

and $\partial_{f, 2} h_{f}$ follows it respectively. The following lemma validates the derivative:
Lemma 10. If the assumptions in Theorem 10 hold, then $\nabla R_{Q^{2}, \lambda}[f]$ is a Gateau-derivative of $R_{Q^{2}, \lambda}$ with the direction $f \in \mathcal{H}_{l}$.

Proof of Lemma 10. We consider a sequence of functions $h_{n} \in \mathcal{H}_{l}$ for $n \in \mathbb{N}$, such that $h_{n}(x) \neq 0, \forall x \in \mathcal{X}$ and $\left\|h_{n}\right\|_{L^{\infty}} \rightarrow 0$ as $n \rightarrow \infty$. Then, for $f \in \mathcal{H}_{l}$, a simple calculation
yields

$$
\begin{aligned}
& \left|\frac{R_{Q^{2}, \lambda}\left(f+h_{n}\right)-R_{Q^{2}, \lambda}(f)-\nabla R_{Q^{2}, \lambda}[f]\left(h_{n}\right)}{\left\|h_{n}\right\|_{L^{\infty}}}\right| \\
& \quad \leq \int\left\|h_{n}\right\|_{L^{\infty}}^{-1} \mid k\left(z, z^{\prime}\right)\left((y-f(x)) h_{n}(x)\right. \\
& \left.\quad+\left(y^{\prime}-f\left(x^{\prime}\right)\right) h_{n}\left(x^{\prime}\right)+h_{n}(x) h_{n}\left(x^{\prime}\right)\right)-\nabla R_{Q^{2}, \lambda}[f]\left(h_{n}\right) \mid d Q^{2}\left(u, u^{\prime}\right) \\
& \leq \int\left\|h_{n}\right\|_{L^{\infty}}^{-1}\left|k\left(z, z^{\prime}\right) h_{n}(x) h_{n}\left(x^{\prime}\right)\right| d Q^{2}\left(u, u^{\prime}\right) \\
& \leq \int\left\|h_{n}\right\|_{L^{\infty}}^{-1}\left|k\left(z, z^{\prime}\right)\left\|h_{n}\right\|_{L^{\infty}}^{2}\right| d Q^{2}\left(u, u^{\prime}\right) \\
& \leq\left\|h_{n}\right\|_{L^{\infty}} \int\left|k\left(z, z^{\prime}\right)\right| d Q^{2}\left(u, u^{\prime}\right) \rightarrow 0,(n \rightarrow \infty) .
\end{aligned}
$$

The convergence follows the definition of $h_{n}$ and the absolute integrability of $k$, which follows the bounded property of $k$ and compactness of $\mathcal{Z}$. Then, we obtain the statement.

Here, we consider its RKHS-type formulation of $\nabla R_{Q^{2}, \lambda}$, which is convenient to describe a minimizer. Let $\Phi_{l}: \mathcal{X} \rightarrow \mathcal{H}_{l}$ be the feature map associated with the RKHS $\mathcal{H}_{l}$, such that $\left\langle\Phi_{l}[x], f\right\rangle_{\mathcal{H}_{l}}=f(x)$ for any $x \in \mathcal{X}$ and $f \in \mathcal{H}_{l}$. Let $\nabla \tilde{R}_{Q^{2}, \lambda}: \mathcal{H}_{l} \rightarrow \mathcal{H}_{l}$ be an operator such that

$$
\nabla \tilde{R}_{Q^{2}, \lambda}(f):=2 \lambda f+\int \partial_{f, 1} h_{f}\left(u, u^{\prime}\right) \Phi_{l}[x](\cdot)+\partial_{f, 2} h_{f}\left(u, u^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot) d Q^{2}\left(u, u^{\prime}\right) .
$$

Obviously, $\nabla R_{Q^{2}, \lambda}[f](\cdot)=\left\langle\nabla \tilde{R}_{Q^{2}, \lambda}(f), \cdot\right\rangle_{\mathcal{H}_{l}}$. Now, we can describe the first-order condition of the minimizer of the risk. Namely, we can state that

$$
\hat{f}=\underset{f \in \mathcal{H}_{l}}{\operatorname{argmin}} R_{Q^{2}, \lambda}(f) \Leftrightarrow \nabla \tilde{R}_{Q^{2}, \lambda}(\hat{f})=0 .
$$

This equivalence follows Theorem 7.4.1 and Lemma 8.7.1 in Luenberger [1997].
Next, we apply the implicit function theorem to obtain an explicit formula of the derivative of $S$. To this end, we consider a second-order derivative $\nabla^{2} \tilde{R}_{Q^{2}, \lambda}: \mathcal{H}_{l} \rightarrow \mathcal{H}_{l}$ as

$$
\nabla^{2} \tilde{R}_{Q^{2}, \lambda}(f):=2 \lambda f+\int k\left(z, z^{\prime}\right)\left(f(x) \Phi_{l}[x](\cdot)+f\left(x^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot)\right) d Q^{2}\left(u, u^{\prime}\right)
$$

which follows (b) in Lemma A. 2 in Hable [2012]. Its basic properties are provided in the following result:

Lemma 11. If Assumption 1 and the assumptions in Theorem 10 hold, then $\nabla^{2} \tilde{R}_{Q^{2}, \lambda}$ is a continuous linear operator and it is invertible.

Proof of Lemma 11. By (b) in Lemma A. 2 in Hable [2012], $\nabla^{2} \tilde{R}_{Q^{2}, \lambda}$ is a continuous linear operator. In the following, we define $A: \mathcal{H}_{l} \rightarrow \mathcal{H}_{l}$ as $A(f)=\int k\left(z, z^{\prime}\right) f(x) \Phi_{l}[x](\cdot)+$
$\left.f\left(x^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot)\right) d Q^{2}\left(u, u^{\prime}\right)$. To show $\nabla^{2} \tilde{R}_{Q^{2}, \lambda}$ is invertible, it is sufficient to show that (i) $\nabla^{2} \tilde{R}_{Q^{2}, \lambda}$ is injective, and (ii) $A$ is a compact operator.

For the injectivity, we fix non-zero $f \in \mathcal{H}_{l}$ and obtain

$$
\begin{aligned}
& \left\|\nabla^{2} \tilde{R}_{Q^{2}, \lambda}(f)\right\|_{\mathcal{H}_{l}}^{2} \\
& =\langle 2 \lambda f+A(f), 2 \lambda f+A(f)\rangle_{\mathcal{H}_{l}} \\
& =4 \lambda^{2}\|f\|_{\mathcal{H}_{l}}^{2}+4 \lambda\langle f, A(f)\rangle_{\mathcal{H}_{l}}+\|A(f)\|_{\mathcal{H}_{l}}^{2} \\
& >4 \lambda\langle f, A(f)\rangle_{\mathcal{H}_{l}} \\
& =4 \lambda\left\langle f, \int k\left(z, z^{\prime}\right) f(x) \Phi[x] d Q^{2}\left(u, u^{\prime}\right)\right\rangle_{\mathcal{H}_{l}}+4 \lambda\left\langle f, \int k\left(z, z^{\prime}\right) f\left(x^{\prime}\right) \Phi\left[x^{\prime}\right] d Q^{2}\left(u, u^{\prime}\right)\right\rangle_{\mathcal{H}_{l}} \\
& =4 \lambda \int k\left(z, z^{\prime}\right) f(x)^{2} d Q^{2}\left(u, u^{\prime}\right)+4 \lambda \int k\left(z, z^{\prime}\right) f\left(x^{\prime}\right)^{2} d Q^{2}\left(u, u^{\prime}\right) \\
& \geq 0 .
\end{aligned}
$$

The last equality follows the property of $\Phi_{l}$ and the last inequality follows the ISPD property in Assumption 1.

For the compactness, we follow Lemma A. 5 in Hable [2012] and obtain that operators $\left(f \mapsto \int k\left(z, z^{\prime}\right) f(x) \Phi_{l}[x](\cdot) d Q^{2}\left(u, u^{\prime}\right)\right)$ and $\left(f \mapsto \int k\left(z, z^{\prime}\right) f\left(x^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot) d Q^{2}\left(u, u^{\prime}\right)\right)$ are compact.

We define the Gateau derivative of $S$. For a functional $F^{\prime} \in \ell^{\infty}\left(\mathcal{G}_{1} \cup \mathcal{G}_{2}\right)$, we define the following function
$\nabla S_{Q^{2}, \lambda}\left(F^{\prime}\right):=-\nabla^{2} \tilde{R}_{Q^{2}, \lambda}^{-1}\left(\int \partial_{f, 1} h_{f_{Q^{2}}}\left(u, u^{\prime}\right) \Phi_{l}[x](\cdot)+\partial_{f, 2} h_{f_{Q^{2}}}\left(u, u^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot) d t\left(F^{\prime}\right)\left(u, u^{\prime}\right)\right)$,
where $f_{Q^{2}}=S_{\lambda}\left(l^{-1}\left(Q^{2}\right)\right)$ and $Q^{2}$ is a signed measure on $\mathcal{W} \times \mathcal{W}$. Then, we provide the following derivative theorem:

Proposition 3. Suppose the assumptions in Theorem 10 hold. For $F \in B_{S}, F^{\prime} \in \ell^{\infty}\left(\mathcal{G}_{1} \cup \mathcal{G}_{2}\right)$, and $s \in \mathbb{R}$ such that $F+s F^{\prime} \in B_{S}, \nabla S_{l(F), \lambda}\left(F^{\prime}\right)$ is a Gateau-derivative of $S_{\lambda}$, namely,

$$
\lim _{s \rightarrow 0}\left\|\frac{S_{\lambda}\left(F+s F^{\prime}\right)-S_{\lambda}(F)}{s}-\nabla S_{l(F), \lambda}\left(F^{\prime}\right)\right\|_{\mathcal{H}_{l}}=0 .
$$

Proof of Proposition 3. This proof has the following two steps, (i) define a proxy operator $\Gamma$, then (ii) prove the statement by the implicit function theorem.
(i) Define $\Gamma$ : Note that $\iota\left(F^{\prime}\right)$ exists since $F+s F^{\prime} \in B_{S}$ implies $F^{\prime} \in B_{0}$. We define the following operator $\Gamma(s, f, \lambda): \mathcal{H}_{l} \rightarrow \mathcal{H}_{l}$ for $f \in \mathcal{H}_{l}$ :

$$
\begin{aligned}
\Gamma(s, f, \lambda):= & \nabla \tilde{R}_{l(F)+s l\left(F^{\prime}\right), \lambda} \\
= & 2 \lambda f+\int \partial_{f, 1} h_{f}\left(u, u^{\prime}\right) \Phi_{l}[x](\cdot)+\partial_{f, 2} h_{f}\left(u, u^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot) d \iota(F)\left(u, u^{\prime}\right) \\
& +s \int \partial_{f, 1} h_{f}\left(u, u^{\prime}\right) \Phi_{l}[x](\cdot)+\partial_{f, 2} h_{f}\left(u, u^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot) d \iota\left(F^{\prime}\right)\left(u, u^{\prime}\right) .
\end{aligned}
$$

For simple derivatives, Lemma A. 2 in Hable [2012] provides

$$
\nabla_{s} \Gamma(s, f, \lambda)=\int \partial_{f, 1} h_{f}\left(u, u^{\prime}\right) \Phi_{l}[x](\cdot)+\partial_{f, 2} h_{f}\left(u, u^{\prime}\right) \Phi_{l}\left[x^{\prime}\right](\cdot) d \iota\left(F^{\prime}\right)\left(u, u^{\prime}\right)
$$

and

$$
\nabla_{f} \Gamma(s, f, \lambda)=\nabla^{2} \tilde{R}_{l(F)+s l\left(F^{\prime}\right), \lambda} .
$$

(ii) Apply the implicit function theorem: By its definition and the optimal conditions, we have

$$
\Gamma(s, f, \lambda)=0 \Leftrightarrow f=S_{\lambda}\left(\iota(F)+s \iota\left(F^{\prime}\right)\right) .
$$

Also, we obtain

$$
\nabla_{f} \Gamma\left(0, S_{\lambda}(F), \lambda\right)=\nabla^{2} \tilde{R}_{\iota(F), \lambda} .
$$

Then, for each $\lambda>0$, by the implicit function theorem, there exists a smooth map $\varphi_{\lambda}: \mathbb{R} \rightarrow \mathcal{H}_{l}$ such that

$$
\Gamma\left(s, \varphi_{\lambda}(s), \lambda\right)=0, \forall s,
$$

and it satisfies

$$
\nabla_{s} \varphi_{\lambda}(0)=-\left(\nabla_{f} \Gamma\left(0, \varphi_{\lambda}(0), \lambda\right)\right)^{-1}\left(\nabla_{s} \Gamma\left(0, \varphi_{\lambda}(0), \lambda\right)\right)=\nabla S_{l(F), \lambda}\left(F^{\prime}\right) .
$$

Also, we have $\varphi_{\lambda}(s)=S\left(Q^{2}+s \mu^{2}\right)$. Then, we have

$$
\begin{aligned}
& \lim _{s \rightarrow 0}\left\|\frac{S_{\lambda}\left(F+s F^{\prime}\right)-S_{\lambda}\left(F^{\prime}\right)}{s}-\nabla S_{l(F), \lambda}\left(F^{\prime}\right)\right\|_{\mathcal{H}_{l}} \\
& =\lim _{s \rightarrow 0}\left\|\frac{\varphi_{\lambda}(s)-\varphi_{\lambda}(0)}{s}-\nabla_{s} \varphi_{\lambda}(0)\right\|_{\mathcal{H}_{l}}=0 .
\end{aligned}
$$

Then, we obtain the statement.

Now, we are ready to prove Theorem 10 and Theorem 17.

Proof of Theorem 17. As a preparation, we mention that $S_{\lambda}$ is differentiable in the Hadamard sense, which is Gateau differentiable by Proposition 3. Lemma A. 7 and A. 8 in Hable [2012] show that $\nabla S_{l(F), \lambda, l(G)}$ is Hadamard-differentiable for any $\lambda, G$ and $F$.

Then, we apply the functional delta method. As shown in Theorem 18 and Lemma 9, we have

$$
\sqrt{n}\left(\iota^{-1}\left(U_{n}^{2}\right)-\iota^{-1}\left(P^{2}\right)\right) \rightsquigarrow 2 \mathbb{G}_{P 1} .
$$

Hence, we obtain

$$
\sqrt{n}\left(\left(\lambda_{0} / \lambda\right) \iota^{-1}\left(U_{n}^{2}\right)-\iota^{-1}\left(P^{2}\right)\right)=\frac{\lambda_{0}}{\lambda} \sqrt{n}\left(\iota^{-1}\left(U_{n}^{2}\right)-\iota^{-1}\left(P^{2}\right)\right)+\frac{\sqrt{n}\left(\lambda-\lambda_{0}\right)}{\lambda} \rightsquigarrow 2 \mathbb{G}_{P^{1}},
$$

since $\lambda-\lambda_{0}=o\left(n^{-1 / 2}\right)$. Utilizing the result, we can obtain

$$
\sqrt{n}\left(\hat{f}-f_{\lambda_{0}}^{*}\right)=\sqrt{n}\left(S_{\lambda_{0}}\left(\left(\lambda_{0} / \lambda\right) \iota^{-1}\left(U_{n}^{2}\right)\right)-S_{\lambda_{0}}\left(\iota^{-1}\left(P^{2}\right)\right)\right)+o_{P}(1) \rightsquigarrow \nabla S_{P^{2}, \lambda_{0}}\left(2 G_{P^{1}}\right),
$$

in $\ell^{\infty}(\mathbb{H})$. The convergence follows the functional delta method.
Proof of Theorem 10. This proof is completed by substituting the Central Limit Theorem part (Theorem 18) of the proof of Theorem 17. From Section 3 in Akritas et al. [1986], the V - and U - processes have the same limit distribution asymptotically, so the same result holds.

## D. 3 Gaussian Process (GP) Interpretation

We present the close connection between the non-parametric model and the Gaussian process (GP) in this section. We will show that the RKHS solution of our objective function (6.6) is equivalent to the maximum function of the posterior distribution of the GP. The relationship is inspired by the similarity of our objective function to that of GLS, and it will be used to derive an efficient cross validation error.

By Mercer theorem, the kernel $l\left(x, x^{\prime}\right)$ can be expanded as $l\left(x, x^{\prime}\right)=\sum_{i=1}^{N} \lambda_{i} \phi_{i}(x) \phi_{i}\left(x^{\prime}\right)$, where $\phi_{i}(x)$ are orthonormal in $L^{2}\left(\mathbb{R}^{d}\right)$, a space of square-integrable real-valued functions, and $N<\infty$ for degenerate kernels and $N=\infty$ otherwise. To satisfy for arbitrary $f(x)=\sum_{i=1}^{N} \alpha_{i} \phi_{i}(x)$ the reproducing property

$$
\left\langle\sum_{i} \alpha_{i} \phi_{i} l(x, \cdot)\right\rangle_{\mathcal{H}_{l}}=f(x),
$$

we choose $\phi_{i}(x)$ such that $\phi_{i}$ is orthogonal in $\mathcal{H}_{l}$ as the Mercer expansion is not unique. Based on the reproducing property, we can obtain $\left\langle\phi_{i}, \phi_{j}\right\rangle_{\mathcal{H}_{l}}=\delta_{i j} \lambda_{i}^{-1}$ where $\delta_{i j}=1$, if $i=j$, and zero otherwise. Similar settings can be found in Walder and Bishop [2017], for example.

Let us consider a GP over a space of functions

$$
f(x)=\Phi(x) \boldsymbol{w}
$$

where $\Phi(x)=\left[\phi_{1}(x), \phi_{2}(x), \ldots, \phi_{N}(x)\right]$ is the feature vector and $w$ is the parameter vector. We assume a prior distribution of $\boldsymbol{w} \sim N(\mathbf{0}, \delta \Lambda)^{1}$ where $\Lambda=\operatorname{diag}\left(\lambda_{1}, \lambda_{2}, \ldots, \lambda_{N}\right)$ is a diagonal matrix with eigenvalues $\lambda_{1}, \ldots, \lambda_{N}$, and $\delta>0$ is a hyper-parameter, which plays the same role as the regularization hyper-parameter (as we show later).

[^4]This definition is equivalent to assuming that on a set of input $x, f(x) \sim \operatorname{GP}(\mathbf{0}, \delta l(x, x))$ because

$$
\begin{aligned}
\mathbb{E}[f(\boldsymbol{x})] & =\Phi(\boldsymbol{x}) \mathbb{E}[\boldsymbol{w}]=\mathbf{0} \\
\operatorname{Var}[f(\boldsymbol{x})] & =\delta \Phi(\boldsymbol{x}) \Lambda \Phi(\boldsymbol{x})^{\top}=\delta l(\boldsymbol{x}, \boldsymbol{x})
\end{aligned}
$$

We refer the interested readers to Rasmussen and Williams [2005] for further details on Gaussian process models.

## D.3.1 Likelihood Function

Given the prior distribution on $w$, we aim to characterize the posterior distribution $p(\boldsymbol{w} \mid D) \propto p(D \mid \boldsymbol{w}) p(\boldsymbol{w})$ where $D=\left\{\left(x_{i}, y_{i}, z_{i}\right)\right\}_{i=1}^{n}$ is an i.i.d. sample of size $n$, and $f(x)=\Phi(x) w$ where the $i$ 's row of $\Phi(x)$ contains the feature vector of $x_{i}$, namely, $\Phi\left(x_{i}\right)^{\top}$. To define the likelihood $p(D \mid \boldsymbol{w})$, we recall from Lemma 1 that the risk $R_{k}(f)$ can be expressed in terms of two independent copies of random variables $(X, Y, Z)$ and $\left(X^{\prime}, Y^{\prime}, Z^{\prime}\right)$. To this end, let $D^{\prime}:=\left\{\left(x_{i}^{\prime}, y_{i}^{\prime}, z_{i}^{\prime}\right)\right\}_{i=1}^{n}$ be an independent sample of size $n$ with an identical distribution to $D$. Given a pair of samples $(x, y, z)$ and $\left(x^{\prime}, y^{\prime}, z^{\prime}\right)$ from $D$ and $D^{\prime}$, respectively, we then define the likelihood as

$$
\begin{aligned}
p\left(\left\{(x, y, z),\left(x^{\prime}, y^{\prime}, z^{\prime}\right)\right\} \mid \boldsymbol{w}\right) \propto & \exp \left[-\frac{1}{2}(y-\Phi(x) \boldsymbol{w}) k\left(z, z^{\prime}\right)\left(y^{\prime}-\Phi\left(x^{\prime}\right) \boldsymbol{w}\right)\right] \\
& =\exp \left[-\frac{1}{2}(y-f(x)) k\left(z, z^{\prime}\right)\left(y^{\prime}-f\left(x^{\prime}\right)\right)\right]
\end{aligned}
$$

Hence, the likelihood on both $D$ and $D^{\prime}$ can be expressed as

$$
\begin{align*}
p\left(\left\{D, D^{\prime}\right\} \mid \boldsymbol{w}\right) \propto & \prod_{i=1}^{n} \prod_{j=1}^{n} p\left(\left\{\left(x_{i}, y_{i}, z_{i}\right),\left(x_{j}^{\prime}, y_{j}^{\prime}, z_{j}^{\prime}\right)\right\} \mid \boldsymbol{w}\right) \\
& =\exp \left[-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right) k\left(z_{i}, z_{j}^{\prime}\right)\left(y_{j}^{\prime}-f\left(x_{j}^{\prime}\right)\right)\right] \tag{D.17}
\end{align*}
$$

In practice, however, we only have access to a single copy of sample, i.e., $D$, but not $D^{\prime}$. One way of constructing $D^{\prime}$ is through data splitting: the original dataset is split into two halves of equal size where the former is used to construct $D$ and the latter is used to form $D^{\prime}$. Unfortunately, this approach reduces the effective sample size that can be used for learning. Alternatively, we propose to estimate the original likelihood (D.17) by using $M$-estimators: given the full dataset $D=\left\{\left(x_{i}, y_{i}, z_{i}\right)\right\}_{i=1}^{n}$, our approximated likelihood can be defined as

$$
\begin{aligned}
p(D \mid \boldsymbol{w}) \propto & \prod_{i=1}^{n} \prod_{j=1}^{n} p\left(\left\{\left(x_{i}, y_{i}, z_{i}\right),\left(x_{j}, y_{j}, z_{j}\right)\right\} \mid \boldsymbol{w}\right) \\
& =\exp \left[-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n}\left(y_{i}-f\left(x_{i}\right)\right) k\left(z_{i}, z_{j}\right)\left(y_{j}-f\left(x_{j}\right)\right)\right]
\end{aligned}
$$

$$
\begin{aligned}
& =\exp \left[-\frac{1}{2}(\boldsymbol{y}-f(\boldsymbol{x}))^{\top} K_{z}(\boldsymbol{y}-f(\boldsymbol{x}))\right] \\
& =(2 \pi)^{n / 2}\left|K_{z}\right|^{-1 / 2} N\left(f(\boldsymbol{x}) \mid \boldsymbol{y}, K_{z}^{-1}\right)
\end{aligned}
$$

The matrix $K_{z}$, as a kernel matrix defined above, plays a role of correlating the residuals $y_{i}-f\left(x_{i}\right)$. Besides, the standard GP regression has a similar likelihood, which is Gaussian but the covariance matrix is computed on $x$ rather than $z$. Based on the likelihood $p(D \mid \boldsymbol{w})$, the maximum likelihood (ML) estimator hence coincides with the minimizer of the non-regularized version of our objective (6.6).

## D.3.2 Maximum A Posteriori (MAP)

Combining the above likelihood function with the prior on $w$, the maximum a posteriori (MAP) estimate of $w$ can be obtained by solving

$$
\begin{align*}
& \underset{w}{\operatorname{argmax}} \log p(w \mid D) \\
&=\underset{w}{\operatorname{argmax}} {[\log p(D \mid \boldsymbol{w})+\log p(\boldsymbol{w})-\log p(D)] } \\
&=\underset{w}{\operatorname{argmax}}[ {\left[-\frac{1}{2}(\boldsymbol{y}-\Phi(\boldsymbol{x}) \boldsymbol{w})^{\top} K_{z}(\boldsymbol{y}-\Phi(\boldsymbol{x}) \boldsymbol{w})\right.} \\
&\left.-\frac{n}{2} \log 2 \pi-\frac{1}{2} \log |\delta \Lambda|-\frac{1}{2} \boldsymbol{w}^{\top}(\delta \Lambda)^{-1} \boldsymbol{w}-\log p(D)\right] \\
&=\underset{w}{\operatorname{argmax}} {\left[-\frac{1}{2}(\boldsymbol{y}-\Phi(\boldsymbol{x}) \boldsymbol{w})^{\top} K_{z}(\boldsymbol{y}-\Phi(\boldsymbol{x}) \boldsymbol{w})-\frac{1}{2} \boldsymbol{w}^{\top}(\delta \Lambda)^{-1} \boldsymbol{w}\right] } \tag{D.18}
\end{align*}
$$

Specifically, setting the first derivative of (D.18) to zero yields the first-order condition

$$
-(\delta \Lambda)^{-1} \hat{w}-\Phi(x)^{\top} K_{z} \Phi(x) \hat{w}+\Phi(x)^{\top} K_{z} y=0
$$

and the MAP estimate

$$
\begin{equation*}
\hat{w}=\delta \Lambda \Phi(x)^{\top} K_{z}\left(y-\Phi(x)^{\top} \hat{w}\right) \tag{D.19}
\end{equation*}
$$

Note that we express $\hat{\boldsymbol{v}}$ using the implicit expression. Given a test data point $x^{*}$, the corresponding prediction can be obtained by computing the posterior expectation, i.e.,

$$
\begin{aligned}
\hat{f}\left(x^{*}\right) & \equiv \mathbb{E}_{f \mid D}\left[f\left(x^{*}\right)\right] \\
& =\Phi\left(x^{*}\right) \hat{w} \\
& =\Phi\left(x^{*}\right)(\delta \Lambda) \Phi(x)^{\top} K_{z}\left(y-\Phi(x)^{\top} \hat{\boldsymbol{w}}\right) \\
& \equiv l\left(x^{*}, \boldsymbol{x}\right) \hat{\boldsymbol{\alpha}}
\end{aligned}
$$

where we define $\hat{\boldsymbol{\alpha}} \equiv \delta K_{z}\left(\boldsymbol{y}-\Phi(\boldsymbol{x})^{\top} \hat{\boldsymbol{w}}\right)$. The second equality holds because $\hat{\boldsymbol{v}}$ is the posterior mean. The third equality is obtained by substituting (D.19) into the second equation. Finally, the last equation follows from the Mercer theorem:
$\Phi\left(x^{\prime}\right) \Lambda \Phi(x)^{\top}=l\left(x^{\prime}, x\right)$.
We can see that without using the representer theorem here, we get the same expression of the optimal solution $\hat{f}$ as the one we obtain by invoking the representer theorem in Section 6.3.2. By substituting $\hat{\boldsymbol{\alpha}}=\delta K_{z}\left(\boldsymbol{y}-\Phi(\boldsymbol{x})^{\top} \hat{\boldsymbol{w}}\right)$ back into (D.19), we obtain

$$
\begin{equation*}
\hat{\boldsymbol{w}}=\Lambda \Phi(\boldsymbol{x})^{\top} \hat{\boldsymbol{\alpha}} . \tag{D.20}
\end{equation*}
$$

Furthermore, putting (D.20) into (D.18) also yields the following optimization problem:

$$
\begin{aligned}
\hat{\boldsymbol{\alpha}} & =\underset{\boldsymbol{\alpha}}{\operatorname{argmin}}\left[\delta^{-1} \boldsymbol{\alpha}^{\top} L \boldsymbol{\alpha}+(\boldsymbol{y}-L \boldsymbol{\alpha})^{\top} K_{z}(\boldsymbol{y}-L \boldsymbol{\alpha})\right] \\
& =\underset{\boldsymbol{\alpha}}{\operatorname{argmin}}\left[\left(n^{2} \delta\right)^{-1} \boldsymbol{\alpha}^{\top} L \boldsymbol{\alpha}+(\boldsymbol{y}-L \boldsymbol{\alpha})^{\top} W(\boldsymbol{y}-L \boldsymbol{\alpha})\right],
\end{aligned}
$$

where $L=l(\boldsymbol{x}, \boldsymbol{x})$ is the kernel matrix. This problem has the same form as that of the $V$-statistic objective (6.7) except that $\left(n^{2} \delta\right)^{-1}$ replaces $\lambda$. Therefore, we conclude that the predictive mean is equivalent to the optimal function of the V-statistic objective (6.7) given that $\left(n^{2} \delta\right)^{-1}=\lambda$, and $\delta$ plays the role of the regularization parameter. Such a GP interpretation is used for an elegant derivation of the efficient cross validation error in Section 6.4.

## D. 4 Related Works in Reinforcement Learning

Firstly, the idea of "kernel loss" was proposed in Feng et al. [2019] to estimate the value function in reinforcement learning. Similar idea has been used to estimate the importance ratio of two state or state-action distributions in Liu et al. [2018] and Uehara et al. [2020]. Kallus [2018] and Kallus [2020] considered the estimation of the average policy effect (APE) and policy learning. Secondly, in the area of causal inference, Wong and Chan [2018] employed similar technique to estimate an average treatment effect. Despite the methodological similarity to our work, these works did not consider the IV regression setting. To better understand the connection, we first introduce the objective functions of the aforementioned works and then highlight the differences, challenges, and novelties of our work.

Kernel loss in reinforcement learning. Although estimating different subjects $f$, Liu et al. [2018] and Feng et al. [2019] employ similar population objective functions in the following form:

$$
\begin{equation*}
\min _{f} \mathbb{E}_{X} \mathbb{E}_{X^{\prime}}\left[((\mathcal{A} f)(X)-f(X)) k\left(X, X^{\prime}\right)\left((\mathcal{A} f)\left(X^{\prime}\right)-f\left(X^{\prime}\right)\right)\right] \tag{D.21}
\end{equation*}
$$

where $\mathcal{A}$ is a task-specific operator acting on $f$ and $X^{\prime}$ is an independent copy of $X$. More specifically, Liu et al. [2018] aims to estimate the importance ratio of state distributions of two policies as $f(X)$. In the task of value function estimation by Feng et al. [2019], $f(X)$ is the value function. In both works, $X$ represents the state variable
in the context of reinforcement learning.

Loss for APE and policy learning. As mentioned by Kallus [2018], the estimation of the APE is similar to that of the average treatment effect, so we introduce the former as an example. By Kallus [2018], the target is to estimate the average effect of a provided policy $\pi_{T}(X)$ for a treatment $T \in[1, \ldots, m]$ conditioned on a covariate $X$. More specifically, given a set of observations of $T, X$ and the outcome $Y,\left\{t_{i}, x_{i}, y_{i}\right\}_{i=1}^{n}$, the APE estimate is $f:=\sum_{i=1}^{n} w_{i} y_{i}$ where $\left\{w_{i}\right\}_{i=1}^{n}$ are unknown weights. We use bold character $z:=\left[z_{i}\right]_{i=1}^{n}$ to denote a column vector of variables $z_{i}$. Then, the weights $\boldsymbol{w}$ are determined by the following objective:

$$
\begin{align*}
& \min _{\boldsymbol{w}}\left(\sum_{T} \gamma_{T}^{q} \mathcal{B}_{T}^{q}\right)^{2 / q}+n^{-2} \boldsymbol{w}^{\top} \Lambda \boldsymbol{w}  \tag{D.22}\\
& \mathcal{B}_{T}(\boldsymbol{w}):=\left(\boldsymbol{w}^{\top} \delta_{t T}-\pi_{T}(\boldsymbol{x})\right)^{\top} k_{T}(\boldsymbol{x}, \boldsymbol{x})\left(\boldsymbol{w}^{\top} \delta_{\boldsymbol{t} T}-\pi_{T}(\boldsymbol{x})\right)
\end{align*}
$$

where $q \in(0,1)$ and $\gamma_{T}>0$ are constants, $\delta_{s, t}=\mathbb{I}[s=t]$ is the Kronecker delta and $\Lambda$ is a positive definite matrix dependent on $\boldsymbol{y}$. Furthermore, $\pi_{T}(\boldsymbol{x})$ and $\delta_{t T}$ are (column) function values evaluated on $x$ and $t$, and $k_{T}(x, x)$ is the kernel matrix on $x$ conditioned on $T$. The policy learning task assumes that no policy is known and requires optimizing both $w$ and $\pi$ on a slight variant of the above objective.

Comparison between related losses (D.21), (D.22) and MMR loss. Compared with our population and empirical risks (6.3) and (6.6), the kernel loss (D.21) and $\mathcal{B}_{T}(\boldsymbol{w})$ in the loss for $\operatorname{APE}(\mathrm{D} .22)$ have a similar quadratic form. The difference is that the kernel loss and $\mathcal{B}_{T}(\boldsymbol{w})$ have variables, which appear in the kernel function, involved in the residual, whereas the kernel function in our objective depends on the instrument $Z$ which does not appear in the residual.

Challenges and novelties of MMR. The difference in losses simplifies the estimation in the related work and requires us to perform new analyses. Specifically, the consistence of the estimators in the above related work holds under mild conditions. That is, in the estimation of the value function, minimizing the population objective function to 0 guarantees a solution for the Bellman equation, which is consistent to the value function according to the unique solution property of the Bellman equation [Feng et al., 2019]; for the importance ratio and the APE estimation, the consistency holds under mild conditions of data distributions [Liu et al., 2018; Kallus, 2018]. In contrast, the estimator by minimizing the MMR risk (6.3) to 0 is hardly consistent to the the true $f(X)$ under mild conditions, because there can be $\hat{f} \neq f$ satisfying $\mathbb{E}[Y-\hat{f}(X) \mid Z]=0$ almost surely. Therefore, we first introduce the completeness condition in Assumption 2 for the estimation, which is classical in the IV regression field and is needed to identify $f(X)$ from the CMR. Second, a new theoretical analysis is necessary to clarify the nature of the estimation on $f(X)$, such as consistency and asymptotic normality. We develop a novel theory for this point in Section 6.5, which has not been studied by these related work.

```
Algorithm 3 MMR-IV (Nyström)
Input: Dataset \(D=\{(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})\}_{i=1}^{n}\), kernel functions \(k\) and \(l\) with parameters \(\theta_{k}\) and
    \(\theta_{l}\), regularization parameter \(\lambda\), leave out data size \(M\), Nyström approximation
    sample size \(M^{\prime}\), LMOCV times \(m\)
Output: predictive value at \(\boldsymbol{x}_{*}\)
    \(K=k\left(\boldsymbol{z}, \boldsymbol{z} ; \theta_{k}\right)\)
    \(\hat{\delta}, \hat{\theta}_{l}=\operatorname{argmin}_{\delta, \theta_{l}} \sum_{i=1}^{m}\left(\boldsymbol{b}^{(i)}-\boldsymbol{y}_{M}^{(i)}\right)^{\top} K_{M}^{(i)}\left(\boldsymbol{b}^{(i)}-\boldsymbol{y}_{M}^{(i)}\right)\)
    \(K / n^{2} \approx \tilde{U} \tilde{V} \widetilde{U}^{\top} \quad\) Nyström Approx with \(M^{\prime}\) samples
    \(\hat{\alpha}=\lambda^{-1}\left[I-\widetilde{u}\left(\lambda^{-1} \tilde{U}^{\top} L \widetilde{U}+\tilde{V}^{-1}\right)^{-1} \tilde{U}^{\top} \lambda^{-1} L\right] \tilde{U} \tilde{v} \tilde{U}^{\top} y\)
    return \(\hat{f}\left(\boldsymbol{x}_{*}\right)=l\left(\boldsymbol{x}_{*}, \boldsymbol{x} ; \hat{\theta}_{l}\right) \hat{\alpha}\)
```

```
Algorithm 4 MMR-IV (NN)
Input: Dataset \(D=\{x, y, z\}\), kernel function \(k\) with parameters \(\theta_{k}, \mathrm{NN} f_{\mathrm{NN}}\) with
    parameters \(\theta_{\mathrm{NN}}\), regularization parameter \(\lambda\)
Output: predictive value at \(\boldsymbol{x}_{*}\)
    Compute \(K=k\left(\boldsymbol{z}, \boldsymbol{z} ; \theta_{k}\right)\)
    \(\hat{f}_{N N}=\operatorname{argmin}_{f_{\mathrm{NN}}}\left(y-f_{\mathrm{NN}}(\boldsymbol{x})\right)^{\top} K\left(y-f_{\mathrm{NN}}(\boldsymbol{x})\right) / n^{2}+\lambda\left\|\theta_{\mathrm{NN}}\right\|_{2}^{2}\)
    return \(\hat{f}_{N N}\left(\boldsymbol{x}_{*}\right)\)
```


## D. 5 Additional Details of Experiments

In this section, we provide additional details about the experiments as well as more discussions on the experimental results presented in Section 6.6.

## D.5.1 Baseline Algorithms

The details of MMR-IV (Nyström) and MMR-IV (NN) algorithms are given in Algorithm 3 and Algorithm 4, respectively. In the experiments, we compare our algorithms to the following baseline algorithms:

- DirectNN: A standard least square regression on $X$ and $Y$ using a neural network (NN).
- 2SLS: A vanilla 2SLS on raw $X$ and $Z$.
- Poly2SLS: The 2SLS that is performed on polynomial features of $X$ and $Z$ via ridge regressions.
- DeepIV [Hartford et al., 2017]: A nonlinear extension of 2SLS using deep NNs. We use the implementation available at https://github.com/microsoft/EconML.
- KernelIV [Singh et al., 2019]: A generalization of 2SLS by modeling relations among $X, Y$, and $Z$ as nonlinear functions in RKHSs. We use the publicly available implementation at https://github.com/r4hu1-5in9h/KIV.
- GMM+NN: An optimally-weighted GMM [Hansen, 1982] is combined with a NN $f(X)$. The details of this algorithm can be found in Bennett et al. [2019, Section 5].
- AGMM [Lewis and Syrgkanis, 2018]: This algorithm models $h(Z)$ by a deep NN and employs a minimax optimization to solve for $f(X)$. The implementation we use is available at https://github.com/vsyrgkanis/adversarial_gmm.
- DeepGMM [Bennett et al., 2019]: This algorithm is a variant of AGMM with optimal inverse-covariance weighting matrix. The publicly available implementation at https://github.com/CausalML/DeepGMM is used and all of the above baselines are provided in the package except KernelIV.
- AGMM-K [Dikkala et al., 2020]: This algorithm extends AGMM by modeling $h(Z)$ and $f(X)$ as RKHSs. Nyström approximation is applied for fast computation. The publicly available implementation at https://github.com/microsoft/Adversa rialGMM is used.


## D.5.2 Experimental Settings on Low-dimensional Scenario

For the experiments in Section 6.6.1, we consider both small-sample ( $n=200$ ) and large-sample ( $n=2000$ ) regimes, in which $n$ points are sampled for training, validation and test sets, respectively. In both regimes, we standardize the values of $Y$ to have zero mean and unit variance for numerical stability. Hyper-parameters of NNs in Algorithm 4, including the learning rate and the regularization parameter, are chosen by 2 -fold CV for fair comparisons with the baselines. These hyper-parameters are provided in Appendix D.5.4. Besides, we use the well-tuned hyper-parameter selections of baselines provided in their packages without changes. We fix the random seed to 527 for all data generation and model initialization.

In contrast to our NN-based method, the RKHS-based method in Algorithm 3 has the analytic form of CV error. We combine the training and validation sets to perform leave-2-out CV to select parameters of the kernel $l$ and the regularization parameter $\lambda$. We choose the sum of Gaussian kernels for $k$ and the Gaussian kernel for $l$. For the Nyström approximation, we subsample 300 points from the combined set. As a small subset of the Gram matrix is used as Nyström samples, which misses much information of data, we avoid outliers in test results by averaging 10 test errors with different Nyström approximations in each experiment. All methods are repeated 10 times on each dataset with random initialization.

Detailed comments on the experimental results. First, under the influence of confounders, DirectNN performs worst as it does not use instruments. Second, MMR - IVs perform reasonably well in both small-sample and large-sample regimes. For the linear function, 2SLS and Poly2SLS tend to outperform other algorithms as the linearity assumption is satisfied in this case. For non-linear functions, some NN based methods show competitive performance in certain cases. Notably, GMM+NN has unstable performance as the function $h$ in (6.1) is designed manually. KernelIV performs quite

Table D.1: The mean square error (MSE) $\pm$ one standard deviation in the small-sample regime ( $n=200$ ).

| Algorithm | abs | linear | sin | step |
| :--- | :---: | :---: | :---: | :---: |
|  | $.143 \pm .000$ | $.046 \pm .000$ | $.404 \pm .006$ | $.253 \pm .000$ |
| 2SLS | $.564 \pm .000$ | $.003 \pm .000$ | $.304 \pm .000$ | $.076 \pm .000$ |
| Poly2SLS | $.125 \pm .000$ | $.003 \pm .000$ | $.164 \pm .000$ | $.077 \pm .000$ |
| GMM+NN | $.792 \pm .000$ | $.203 \pm .000$ | $1.56 \pm .001$ | $.550 \pm .000$ |
| AGMM | $.031 \pm .000$ | $.011 \pm .000$ | $.330 \pm .000$ | $.080 \pm .000$ |
| DeepIV | $.204 \pm .008$ | $.047 \pm .004$ | $.197 \pm .004$ | $.039 \pm .001$ |
| DeepGMM | $.022 \pm .003$ | $.032 \pm .016$ | $.143 \pm .030$ | $.039 \pm .002$ |
| KernelIV | $.063 \pm .000$ | $.024 \pm .000$ | $.086 \pm .000$ | $.055 \pm .000$ |
| AGMM-K | $12.3 \pm .000$ | $1.32 \pm .000$ | $1.57 \pm .000$ | $1.71 \pm .000$ |
| DualIV | $.202 \pm .000$ | $0.103 \pm .000$ | $.251 \pm .000$ | $.362 \pm .000$ |
| MMR-IV (NN) | $.019 \pm .003$ | $.004 \pm .001$ | $.292 \pm .024$ | $.075 \pm .008$ |
| MMR-IV (RKHS) | $.030 \pm .000$ | $.011 \pm .000$ | $.075 \pm .000$ | $.057 \pm .000$ |

well but not as well as our method. AGMM-K is similar in principle to our method while its errors are high. We suspect that this is due to the hyper-parameter selection is not flexible enough. Additionally, we observe that the AGMM-K's performance becomes better as the numbers of CV folds and of hyper-parameter candidates increase. A similar observation is also obtained on DualIV. Thus, it show that it is desirable to have the analytical CV error, which is a advantage of our method against other RKHS baselines. Besides, the selection of the weight matrix remains an open question, and although DualIV, AGMM-K and MMRIV (RKHS) rely on the median heuristic to select the bandwidth, we believe that the selection has different effects on different methods and it is difficult to get fair selection. We will leave this problem to future work. Moreover, the performances of the complicated methods like DeepIV and DeepGMM deteriorate more in the large-sample regimes than the small-sample regimes. We suspect that this is because these methods rely on two NNs and are thus sensitive to different hyper-parameters. In contrast, MMR-IV (Nyström) has the advantage of adaptive hyper-parameter selection.

## D.5.3 Experimental Settings on High-dimensional Scenario

For experiments in Section 6.6.2, we sample $n=10,000$ points for the training, validation, and test sets, and run each method 10 times. For MMR-IV (Nyström), we run it only on the training set to reduce computational workload, and use principal component analysis (PCA) to reduce dimensions of $X$ from 728 to 8 . We still use the sum of Gaussian kernels for $k\left(z, z^{\prime}\right)$, but use the automatic relevance determination (ARD) kernel for $l\left(x, x^{\prime}\right)$. We omit KernelIV and DualIV in this experiment since their kernel parameter selection is not suitable for image data.

Detailed comments on the experimental results. Like Bennett et al. [2019], we
observe that the DeepIV code returns NaN when $X$ is high-dimensional. Besides, we run Ridge2SLS, which is Poly2SLS with fixed linear degree, in place of Poly2SLS to reduce computational workload. AGMM-K+NN, i.e., the KLayerTrained method in [Dikkala et al., 2020], is a variant of AGMM-K, which uses neural networks to model $f(X)$, and to extract features of $Z$ as inputs of the kernel $k$. To reinforce the flexibility, AGMM $-K+N N$ also trains the kernel hyper-parameter, i.e., the lengthscale of the RBF kernel $k$. 2SLS has large errors for high-dimensional $X$ because the first-stage regression from $Z$ to $X$ is ill-posed. The average performance of GMM + NN suggests that manually designed functions for instruments is insufficient to extract useful information. Furthermore, MMR-IV (NN) performs competitively across scenarios. On MNIST $_{\text {Z }}$, MMR - IVs perform better than other methods, which implies using the sum of Gaussian kernels for the kernel $k$ is proper. DeepGMM has competitive performance as well. On MNIST $X_{X}$ and MNIST $_{X Z}$, MMR-IV (NN) outperforms all other methods. Compared with MMR-IV (NN), AGMM-K+NN has a more flexible kernel $k$ but is lack of good kernel hyper-parameter selection. So, it is vulnerable to local minimum and shows unstable performance. The ARD kernel with PCA of MMR-IV (Nyström) fails because the features from PCA are not representative enough. In addition, we observe that DeepGMM often produces results that are unreliable across all settings. We suspect that hard-to-optimize objective and complicated optimization procedure are the causes. Compared with DirectNN, most of baselines can hardly deal with high-dimensional structured instrumental regressions.

## D.5.4 More Details of MMR-IVs

For the kernel function on instruments, we employ the sum of Gaussian kernels

$$
k\left(z, z^{\prime}\right)=\frac{1}{3} \sum_{i=1}^{3} \exp \left(-\frac{\left\|z-z^{\prime}\right\|_{2}^{2}}{2 \sigma_{k i}^{2}}\right)
$$

and the Gaussian kernel for $l\left(x, x^{\prime}\right)=\exp \left(-\left\|x-x^{\prime}\right\|_{2}^{2} /\left(2 \sigma_{l}^{2}\right)\right)$, where $\sigma_{k 1}$ is chosen as the median interpoint distance of $z=\left\{z_{i}\right\}_{i=1}^{n}$ and $\sigma_{k 2}=0.1 \sigma_{k 1}, \sigma_{k 3}=10 \sigma_{k 1}$. The motivation of such a kernel $k$ is to optimize $f$ on multiple kernels, and we leave parameter selection of $k$ to the future work.

As per the dimensions of $X$, we parametrize $f$ either as a fully connected neural network with leaky ReLU activations and 2 hidden layers, each of which has 100 cells, for non-image data, or a deep convolutional neural network (CNN) architecture for MNIST data. We denote the fully connected neural network as $\operatorname{FCNN}(100,100)$ and refer readers to our code release for exact details on our CNN construction. Learning rates and regularization parameters are summarized in Table D.2.

## D.5.5 Additional Comments on Mendelian Randomization

For experiments in Section 6.6.3, KernelIV also achieves competitive and stable performance across all settings, but not as good as ours. DirectNN is always among

Table D.2: Hyper-parameters of neural networks used in the experiments.

| Scenario | Model $f$ | Learning Rates | $\lambda$ |
| :---: | :---: | :---: | :---: | :---: |
| Low Dimensional | FCNN $(100,100)$ | $\left(10^{-12}, 10^{-11}, 10^{-10}, 10^{-9}, 10^{-8}, 10^{-7}, 10^{-6}\right)$ | $\left(5 \times 10^{-5}, 10^{-4}, 2 \times 10^{-4}\right)$ |
| MNIST $_{Z}$ | FCNN $(100,100)$ | $\left(10^{-12}, 10^{-11}, 10^{-10}, 10^{-9}, 10^{-8}, 10^{-7}, 10^{-6}\right)$ | $\left(5 \times 10^{-5}, 10^{-4}, 2 \times 10^{-4}\right)$ |
| MNIST $_{X}$ | CNN | $\left(10^{-12}, 10^{-11}, 10^{-10}, 10^{-9}, 10^{-8}, 10^{-7}, 10^{-6}\right)$ | $\left(5 \times 10^{-5}, 10^{-4}, 2 \times 10^{-4}\right)$ |
| MNIST $_{\text {XZ }}$ | CNN | $\left(10^{-12}, 10^{-11}, 10^{-10}, 10^{-9}, 10^{-8}, 10^{-7}, 10^{-6}\right)$ | $\left(5 \times 10^{-5}, 10^{-4}, 2 \times 10^{-4}\right)$ |
| Mendelian | FCNN $(100,100)$ | $\left(10^{-12}, 10^{-11}, 10^{-10}, 10^{-9}, 10^{-8}, 10^{-7}, 10^{-6}\right)$ | $\left(5 \times 10^{-5}, 10^{-4}, 2 \times 10^{-4}\right)$ |

the worst approaches on all settings as no instrument is used. Poly2SLS performs accurately on the last two experiments, while presents significant instability with the number of instruments in Figure 6.2(left) because of failure of the hyper-parameter selection. In Figure 6.2 (middle) and Figure 6.2(right), we can observe that the performance of most approaches deteriorates as the effect of confounders becomes stronger. MMR-IV (Nyström) has promising performance and shows a bit more sensitivity to $c_{2}$ than $c_{1}$, and the good performance take the advantage of the hyper-parameter selection compared with AGMM-K and DualIV.

## D.5.6 Experimental Settings on Vitamin D Data

We normalize each variable to have a zero mean and unit variance to reduce the influence of different scales. We consider two cases: (i) without instrument and (ii) with instruments. By without instrument, we mean that the $W_{V}$ matrix of MMR-IV (Nyström) becomes an identity matrix. Following Sjolander and Martinussen [2019], we assess the effect of Vitamin D (exposure) on mortality rate (outcome), control the age in the analyses, and use filaggrin as the instrument. We illustrate original Vitamin D, age and death in Figure D.2. We randomly pick (random seed is 527) 300 Nyström samples and use leave-2-out cross validation to select hyper-parameters. The generalized linear models in [Sjolander and Martinussen, 2019] are a linear function in the first step and a logistic regression model in the second step.

In this experiment, we use age as a control variable by considering a structural equation model, which is similar to the model (2.13) except the presence of the controlled (exogenous) variable $C$,

$$
Y=f(X, C)+\varepsilon, \quad X=t(Z, C)+g(\varepsilon)+v
$$

where $\mathbb{E}[\varepsilon]=0$ and $\mathbb{E}[v]=0$. We further assume that the instrument $Z$ satisfies the following three conditions:
(i) Relevance: $Z$ has a causal influence on $X$;
(ii) Exclusion restriction: $Z$ affects $Y$ only through $X$, i.e., $Y \Perp Z \mid X, \varepsilon, C$;
(iii) Unconfounded instrument(s): Z is conditionally independent of the error, i.e., $\varepsilon \Perp Z \mid C$.

Unlike the conditions specified in the main text, (ii) and (iii) also include the controlled variable C. A similar model is employed in Hartford et al. [2017]. From Assumption
(iii), we can see that $\mathbb{E}[\varepsilon \mid C, Z]=\mathbb{E}[\varepsilon \mid C]$, and based on this, we further obtain

$$
\mathbb{E}[(Y-f(X, C)-\mathbb{E}[\varepsilon \mid C]) h(Z, C)]=0
$$

for any measurable function $h$. Note that $\mathbb{E}[\varepsilon \mid C]$ is only conditioned on $C$, remains constant on arbitrary values of $C$, and is typically non-zero. To adapt our method to this model, we only need to use the kernel $k$ with $(Z, C)$ as inputs and be aware that the output of the method is an estimate of $f^{\prime}(X, C):=f(X, C)+\mathbb{E}[\varepsilon \mid C]$ instead of just $f(X, C)$.


Figure D.1: Estimated effect of vitamin D level on mortality rate, controlled by age. All plots depict normalized contours of $\hat{f}^{\prime}$ defined in Section D.5.6 where blue represents low mortality rate and yellow the opposite. We can divide each plot roughly into left (young age) and right (old age) parts. While the right parts reflect similar information (i.e., lower vitamin D level at an old age leads to higher mortality rate), the left parts are different. In (a), a high level of vitamin D at a young age can result in a high mortality rate, which is counter-intuitive. A plausible explanation is that it is caused by some unobserved confounders between vitamin D level and mortality rate. In (b), on the other hand, this spurious effect disappears when the filaggrin mutation is used as instrument, i.e., a low vitamin $D$ level at a young age has only a slight effect on death, but a more adverse effect at an old age [Meehan and Penckofer, 2014]. This comparison demonstrates the benefit of an instrument variable. (c) and (d) correspond to the results obtained by using Sjolander and Martinussen [2019]' generalized linear model (GLM), from which we can draw similar conclusions. It is noteworthy that MMRIV allows more flexible non-linearity for causal effect.


Figure D.2: Distribution of Vitamin D data. Data points are plotted in the middle, the solid curve and histogram on the right describe the kernel density estimation and histogram of Vitamin D, and those on the top are for Age.

## Bibliography

Achab, M.; Bacry, E.; Gaïffas, S.; Mastromatteo, I.; and Muzy, J.-F., 2017. Uncovering causality from multivariate hawkes integrated cumulants. J. Mach. Learn. Res., 18, 1 (Jan. 2017), 6998-7025. (cited on page 18)

Adams, R. P.; Murray, I.; and MacKay, D. J. C., 2009. Tractable nonparametric bayesian inference in poisson processes with gaussian process intensities. International Conference on Machine Learning (ICML), (2009), 9-16. (cited on page 24)

Ai, C. and Chen, X., 2003. Efficient estimation of models with conditional moment restrictions containing unknown functions. Econometrica, 71, 6 (2003), 1795-1843. (cited on pages 15 and 20)

Akritas, M. G. et al., 1986. Empirical processes associated with V-statistics and a class of estimators under random censoring. The Annals of Statistics, 14, 2 (1986), 619-637. (cited on page 115)

Ambrogioni, L.; Güçlü, U.; GüçLütürk, Y.; Hinne, M.; van Gerven, M. A.; and Maris, E., 2018. Wasserstein variational inference. In Advances in Neural Information Processing Systems, 2473-2482. (cited on pages 19 and 51)

Ancarani, L. U. and Gasaneo, G., 2008. Derivatives of any order of the confluent hypergeometric function $\mathrm{f} 11(\mathrm{a}, \mathrm{b}, \mathrm{z})$ with respect to the parameter a or b . Journal of Mathematical Physics, 49, 6 (2008), 063508. (cited on page 39)

Andersen, T. G. and Sorensen, B. E., 1996. GMM estimation of a stochastic volatility model: A monte carlo study. Journal of Business \& Economic Statistics, 14, 3 (1996), 328-352. (cited on page 61)

Angrist, J. D.; Imbens, G. W.; and Rubin, D. B., 1996. Identification of causal effects using instrumental variables. Journal of the American Statistical Association, 91, 434 (1996), 444-455. (cited on pages 21 and 22)

Angrist, J. D. and Pischke, J.-S., 2008. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press. (cited on pages 60 and 71)

Arcones, M. A. and Gine, E., 1993. Limit theorems for U-processes. The Annals of Probability, (1993), 1494-1542. (cited on pages 108 and 109)

Arjovsky, M.; Chintala, S.; and Bottou, L., 2017. Wasserstein gan. arXiv preprint arXiv:1701.07875, (2017). (cited on pages 19 and 51)

Aronszajn, N., 1950. Theory of reproducing kernels. Transactions of the American Mathematical Society, 68, 3 (1950), 337-404. (cited on page 61)

Athey, S.; Tibshirani, J.; Wager, S.; et al., 2019. Generalized random forests. Annals of Statistics, 47, 2 (2019), 1148-1178. (cited on page 20)

Attias, H., 1999. Inferring parameters and structure of latent variable models by variational bayes. In Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence, 21-30. Morgan Kaufmann Publishers Inc. (cited on page 40)

Bacry, E.; Bompaire, M.; Gaïffas, S.; and Poulsen, S., 2017. tick: a Python library for statistical learning, with a particular emphasis on time-dependent modeling. ArXiv e-prints, (2017). (cited on pages 29 and 43)

Bacry, E.; Mastromatteo, I.; and Muzy, J.-F., 2015. Hawkes processes in finance. Market Microstructure and Liquidity, 1, 01 (2015), 1550005. (cited on page 1)

Bacry, E. and Muzy, J.-F., 2016. First- and second-order statistics characterization of hawkes processes and non-parametric estimation. IEEE Transactions on Information Theory, 62, 4 (2016), 2184-2202. (cited on pages $2,4,17,18,29$, and 45)

Bao, P.; Shen, H.-W.; Jin, X.; and Cheng, X.-Q., 2015. Modeling and predicting popularity dynamics of microblogs using self-excited hawkes processes. In Proceedings of the 24th International Conference on World Wide Web, 9-10. ACM. (cited on page 2)

Barabási, A.-L., 2005. The origin of bursts and heavy tails in human dynamics. Nature, 435, 7039 (may 2005), 207-11. doi:10.1038/nature03459. URL http: //dx.doi.org/10.1038/nature03459. (cited on page 2)

Benamou, J.-D. and Brenier, Y., 2000. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. Numerische Mathematik, 84, 3 (Jan 2000), 375-393. doi:10.1007/s002110050002. (cited on page 20)

Bennett, A.; Kallus, N.; and Schnabel, T., 2019. Deep generalized method of moments for instrumental variable analysis. In Advances in Neural Information Processing Systems 32, 3564-3574. Curran Associates, Inc. (cited on pages 16, 20, 21, $22,61,68,70,121$, and 122)

Berlinet, A. and Thomas-Agnan, C., 2004. Reproducing Kernel Hilbert Spaces in Probability and Statistics. Kluwer Academic Publishers. (cited on page 61)

Bishor, C. M., 2006. Pattern Recognition and Machine Learning. Springer. (cited on pages 11 and 12)

Blundell, R.; Chen, X.; and Kristensen, D., 2007. Semi-nonparametric IV estimation of shape-invariant Engel curves. Econometrica, 75, 6 (2007), 1613-1669. (cited on page 21)

Bonneel, N.; Rabin, J.; Peyré, G.; and Pfister, H., 2015. Sliced and radon wasserstein barycenters of measures. Journal of Mathematical Imaging and Vision, 51, 1 (2015), $22-45$. (cited on page 51)

Bui, T. D.; Yan, J.; and Turner, R. E., 2017. A Unifying Framework for Gaussian Process Pseudo-point Approximations Using Power Expectation Propagation. J. Mach. Learn. Res., 18, 1 (Jan. 2017), 3649-3720. (cited on page 19)

Burgess, S.; Foley, C. N.; Allara, E.; Staley, J. R.; and Howson, J. M. M., 2020. A robust and efficient method for Mendelian randomization with hundreds of genetic variants. Nature Communications, 11, 1 (2020), 376. (cited on pages 60 and 70)

Burgess, S.; Small, D. S.; and Thompson, S. G., 2017a. A review of instrumental variable estimators for mendelian randomization. Statistical Methods in Medical Research, 26, 5 (2017), 2333-2355. (cited on page 20)

Burgess, S.; Small, D. S.; and Thompson, S. G., 2017b. A review of instrumental variable estimators for Mendelian randomization. Statistical Methods in Medical Research, 26, 5 (2017), 2333-2355. (cited on page 60)

Byrd, R. H.; Lu, P.; Nocedal, J.; and Zhu, C., 1995. A limited memory algorithm for bound constrained optimization. SIAM Journal on Scientific Computing, 16, 5 (1995), 1190-1208. (cited on page 54)

Card, D., 1999. The causal effect of education on earnings. In Handbook of Labor Economics (Eds. O. Ashenfelter and D. Card), vol. 3 of Handbook of Labor Economics, chap. 30, 1801-1863. Elsevier. (cited on page 60)

Carrasco, M., 2012. A regularization approach to the many instruments problem. Journal of Econometrics, 170, 2 (2012), 383-398. (cited on page 61)

Carrasco, M. and Florens, J.-P., 2000. Generalization of gmm to a continuum of moment conditions. Econometric Theory, 16, 6 (2000), 797-834. (cited on page 61)

Carrasco, M. and Florens, J.-P., 2014. On the asymptotic efficiency of gmm. Econometric Theory, 30, 2 (2014), 372-406. doi:10.1017/S0266466613000340. (cited on page 61)

Carrasco, M.; Florens, J.-P.; and Renault, E., 2007. Linear inverse problems in structural econometrics estimation based on spectral decomposition and regularization. In Handbook of Econometrics (Eds. J. Нeckman and E. Leamer), vol. 6B, chap. 77. Elsevier, 1 edn. (cited on pages 61 and 64)

Celeux, G. and Diebolt, J., 1985. The sem algorithm: A probabilistic teacher algorithm derived from the EM algorithm for the mixture problem. Comp. Stat. Quarterly, 2 (1985), 73-82. (cited on pages 4, 23, and 28)

Chen, X. and Christensen, T. M., 2018. Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric IV regression. Quantitative Economics, 9, 1 (2018), 39-84. (cited on page 21)

Chen, X. and Pouzo, D., 2012. Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. Econometrica, 80, 1 (2012), 277-321. (cited on page 21)

Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; and Robins, J., 2018. Double/debiased machine learning for treatment and structural parameters. (cited on page 20)

Courty, N.; Flamary, R.; Tuia, D.; and Corpetti, T., 2016. Optimal transport for data fusion in remote sensing. In 2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), 3571-3574. IEEE. (cited on page 50)

Cuturi, M., 2013. Sinkhorn distances: Lightspeed computation of optimal transport. In Advances in neural information processing systems, 2292-2300. (cited on page 51)

Daley, D. J. and Vere-Jones, D., 2003. An introduction to the theory of point processes: volume I: Elementary Theory and Methods. Springer Science \& Business Media. (cited on page 7)

Darolles, S.; Fan, Y.; Florens, J. P.; and Renault, E., 2011. Nonparametric instrumental regression. Econometrica, 79, 5 (2011), 1541-1565. (cited on page 21)

Daskalakis, C. and Panageas, I., 2018. The limit points of (optimistic) gradient descent in min-max optimization. In Proceedings of the 32nd International Conference on Neural Information Processing Systems, 9256-9266. (cited on page 21)

Davey Smith, G. and Ebrahim, S., 2003. 'mendelian randomization': can genetic epidemiology contribute to understanding environmental determinants of disease? International journal of epidemiology, 32, 1 (2003), 1-22. (cited on page 20)
de Jong, R. M., 1996. The bierens test under data dependence. Journal of Econometrics, 72, 1 (1996), 1-32. (cited on page 61)

Dehaene, G. and Barthelmé, S., 2018. Expectation propagation in the large data limit. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 80, 1 (2018), 199-217. (cited on page 20)

Dempster, A. P.; Laird, N. M.; and Rubin, D. B., 1977. Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society, Series B, 39, 1 (1977), 1-38. (cited on pages 4, 23, and 28)

Dezfouli, A. and Bonilla, E. V., 2015. Scalable inference for gaussian process models with black-box likelihoods. In Advances in Neural Information Processing Systems, 1414-1422. (cited on page 19)

D'Haultfoeuille, X., 2011. On the completeness condition in nonparametric instrumental problems. Econometric Theory, 27, 3 (2011), 460-471. URL http: //www.jstor.org/stable/27975488. (cited on page 62)

Dieng, A. B.; Tran, D.; Ranganath, R.; Paisley, J.; and Blei, D., 2017. Variational Inference via \chi Upper Bound Minimization. In Advances in Neural Information Processing Systems 30 (Eds. I. Guyon; U. V. Luxburg; S. Bengio; H. Wallach; R. Fergus; S. Vishwanathan; and R. Garnett), 2732-2741. Curran Associates, Inc. (cited on page 50)

Diggle, P. J.; Moraga, P.; Rowlingson, B.; and Taylor, B. M., 2013. Spatial and spatio-temporal log-gaussian cox processes: extending the geostatistical paradigm. Statistical Science, 28, 4 (2013), 542-563. (cited on page 24)

Dikkala, N.; Lewis, G.; Mackey, L.; and Syrgkanis, V., 2020. Minimax estimation of conditional moment models. CoRR, abs/2006.07201 (2020). (cited on pages 20, 22, 68,121 , and 123)

Dominguez, M. and Lobato, I., 2004. Consistent estimation of models defined by conditional moment restrictions. Econometrica, 72, 5 (2004), 1601-1615. (cited on page 61)

Donald, S. G.; Imbens, G. W.; and Newey, W. K., 2003. Empirical likelihood estimation and consistent tests with conditional moment restrictions. Journal of Econometrics, 117, 1 (2003), 55-93. doi:https://doi.org/10.1016/S0304-4076(03)00118-0. URL https://www.sciencedirect.com/science/article/pii/S0304407603001180. (cited on page 61)

Donald, S. G. and Newey, W. K., 2001. Choosing the number of instruments. Econometrica, 69, 5 (2001), 1161-1191. (cited on page 61)

Donnet, S.; Rivoirard, V.; and Rousseau, J., 2018. Nonparametric bayesian estimation of multivariate hawkes processes. arXiv preprint arXiv:1802.05975, (2018). (cited on pages 3,18 , and 19)

Dowson, D. and Landau, B., 1982. The fréchet distance between multivariate normal distributions. Journal of Multivariate Analysis, 12, 3 (1982), 450 - 455. (cited on page 20)

Du, N.; Dai, H.; Trivedi, R.; Upadhyay, U.; Gomez-Rodriguez, M.; and Song, L., 2016. Recurrent marked temporal point processes: Embedding event history to vector. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1555-1564. ACM. (cited on pages 17 and 18)

Dua, D. and Graff, C., 2017. UCI machine learning repository. URL http://archive.ic s.uci.edu/ml. (cited on page 55)

Eichler, M.; Dahlhaus, R.; and Dueck, J., 2017. Graphical modeling for multivariate hawkes processes with nonparametric link functions. Journal of Time Series Analysis, 38, 2 (2017), 225-242. (cited on pages $2,4,17$, and 18)

El Moselhy, T. A. and Marzouk, Y. M., 2012. Bayesian inference with optimal maps. Journal of Computational Physics, 231, 23 (2012), 7815-7850. (cited on page 20)

Feng, Y.; Li, L.; and Liu, Q., 2019. A kernel loss for solving the bellman equation. Advances in neural information processing systems, 32 (2019). (cited on pages 118 and 119)

Filimonov, V. and Sornette, D., 2015. Apparent criticality and calibration issues in the hawkes self-excited point process model: application to high-frequency financial data. Quantitative Finance, 15, 8 (2015), 1293-1314. (cited on page 2)

Flannery, B. P.; Press, W. H.; Teukolsky, S. A.; and Vetterling, W., 1992. Numerical recipes in C. Press Syndicate of the University of Cambridge, New York, 24 (1992), 78. (cited on page 65)

Flaxman, S.; Teh, Y. W.; and Sejdinovic, D., 2017. Poisson intensity estimation with reproducing kernels. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, vol. 54 of Proceedings of Machine Learning Research, 270-279. PMLR, Fort Lauderdale, FL, USA. URL http://proceedings.mlr.press/v54/flaxman17 a.html. (cited on pages $17,24,26,56$, and 87 )

Fukumizu, K.; Bach, F. R.; and Jordan, M. I., 2004. Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. Journal of Machine Learning Research, 5 (Dec. 2004), 73-99. (cited on page 97)

Gelbrich, M., 1990. On a Formula for the $L_{2}$ Wasserstein Metric between Measures on Euclidean and Hilbert Spaces. Mathematische Nachrichten, 147, 1 (1990), 185-203. doi:10.1002/mana.19901470121. (cited on page 20)

Gelman, A.; Vehtari, A.; Jylänki, P.; Sivula, T.; Tran, D.; Sahai, S.; Blomstedt, P.; Cunningham, J. P.; Schiminovich, D.; and Robert, C., 2017. Expectation propagation as a way of life: A framework for bayesian inference on partitioned data. arXiv preprint arXiv:1412.4869, (2017). (cited on pages 3, 20, and 50)

GPY, since 2012. GPy: A Gaussian process framework in python. http://github.com/S heffieldML/GPy. (cited on page 54)

Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. C., 2017. Improved training of wasserstein gans. In Advances in neural information processing systems, 5767-5777. (cited on page 19)

Hable, R., 2012. Asymptotic normality of support vector machine variants and other regularized kernel methods. Journal of Multivariate Analysis, 106 (2012), 92-117. (cited on pages 108, 112, 113, and 114)

Hall, A., 2005. Generalized Method of Moments. Advanced texts in econometrics. Oxford University Press. (cited on pages 61 and 64)

Hall, A. R. et al., 2005. Generalized method of moments. Oxford university press. (cited on page 61)

Hall, P. and Horowitz, J. L., 2005. Nonparametric methods for inference in the presence of instrumental variables. The Annals of Statistics, 33, 6 (12 2005), 2904-2929. (cited on page 21)

Halpin, P. F., 2012. An em algorithm for hawkes process. Psychometrika, 2 (2012). (cited on pages 27 and 42)

Hansen, L. P., 1982. Large sample properties of generalized method of moments estimators. Econometrica: Journal of the Econometric Society, (1982), 1029-1054. (cited on pages 21,61 , and 121)

Hartford, J.; Lewis, G.; Leyton-Brown, K.; and Taddy, M., 2017. Deep IV: A flexible approach for counterfactual prediction. In Proceedings of the 34th International Conference on Machine Learning, vol. 70, 1414-1423. PMLR. (cited on pages 16, 21, $22,68,120$, and 124)

Hartford, J. S.; Veitch, V.; Sridhar, D.; and Leyton-Brown, K., 2020. Valid causal inference with (some) invalid instruments. CoRR, abs/2006.11386 (2020). (cited on page 70)

Hartwig, F. P.; Davey Smith, G.; and Bowden, J., 2017. Robust inference in summary data Mendelian randomization via the zero modal pleiotropy assumption. International journal of epidemiology, 46, 6 (2017), 1985-1998. (cited on page 70)

Hawkes, A. G., 1971. Spectra of some self-exciting and mutually exciting point processes. Biometrika, 58, 1 (1971), 83-90. (cited on pages 1 and 8)

Hawkes, A. G. and Oakes, D., 1974. A cluster process representation of a self-exciting process. Journal of Applied Probability, 11, 3 (1974), 493-503. (cited on page 1)

Heess, N.; Tarlow, D.; and Winn, J., 2013. Learning to pass expectation propagation messages. In Advances in Neural Information Processing Systems, 3219-3227. (cited on pages 5 and 50)

Hensman, J.; Matthews, A.; and Ghahramani, Z., 2015. Scalable variational Gaussian process classification. Journal of Machine Learning Research, (2015). (cited on page 50)

Hensman, J.; Zwiessele, M.; and Lawrence, N., 2014. Tilted variational bayes. In Artificial Intelligence and Statistics, 356-364. (cited on pages 19 and 50)

Hernández-Lobato, J. M.; Li, Y.; Rowland, M.; Hernández-Lobato, D.; Bui, T.; and Turner, R., 2016. Black-box $\alpha$-divergence minimization. International Conference on Machine Learning, (2016). (cited on pages 5 and 19)

Hoeffding, W., 1963. Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association, 58, 301 (1963), 13-30. (cited on page 101)

Horowitz, J. L., 2011. Applied nonparametric instrumental variables estimation. Econometrica, 79, 2 (2011), 347-394. (cited on page 21)

Huang, G.; Guo, C.; Kusner, M. J.; Sun, Y.; Sha, F.; and Weinberger, K. Q., 2016. Supervised word mover's distance. In Advances in Neural Information Processing Systems, 4862-4870. (cited on page 50)

Ishwaran, H. and James, L. F., 2001. Gibbs sampling methods for stick-breaking priors. Journal of the American Statistical Association, 96, 453 (2001), 161-173. (cited on pages 4,23 , and 25 )

Ivanov, V. K.; Vasin, V. V.; and Tanana, V., 2002. Theory of Linear Ill-posed Problems and Its Applications. Inverse and ill-posed problems series. VSP. (cited on page 22)

Jarrett, R., 1979. A note on the intervals between coal-mining disasters. Biometrika, 66, 1 (1979), 191-193. (cited on page 56)

Jordan, M. I.; Ghahramani, Z.; Jaakkola, T. S.; and Saul, L. K., 1999. An introduction to variational methods for graphical models. Machine learning, 37, 2 (1999), 183-233. (cited on pages 12,19 , and 50 )

Jylänki, P.; Vanhatalo, J.; and Vehtari, A., 2011. Robust gaussian process regression with a student-t likelihood. Journal of Machine Learning Research, 12, Nov (2011), 3227-3257. (cited on pages 13 and 50)

Kallus, N., 2018. Balanced policy evaluation and learning. Advances in neural information processing systems, 31 (2018). (cited on pages 118 and 119)

Kallus, N., 2020. Generalized optimal matching methods for causal inference. Journal of Machine Learning Research, 21, 62 (2020), 1-54. (cited on page 118)

Kariya, T. and Kurata, H., 2004. Generalized least squares. John Wiley \& Sons. (cited on page 63)

Klungel, O.; Jamal Uddin, M.; de Boer, A.; Belitser, S.; Groenwold, R.; and Roes, K., 2015. Instrumental variable analysis in epidemiologic studies: an overview of the estimation methods. Pharm Anal Acta, 6, 353 (2015), 2. (cited on page 60)

Kuang, Z.; Sala, F.; Sohoni, N.; Wu, S.; Córdova-Palomera, A.; Dunnmon, J.; Priest, J.; and Re, C., 2020. Ivy: Instrumental variable synthesis for causal inference. In Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, vol. 108 of Proceedings of Machine Learning Research, 398-410. PMLR. (cited on page 70 )

Kurashima, T.; Althoff, T.; and Leskovec, J., 2018. Modeling interdependent and periodic real-world action sequences. In Proceedings of the 2018 World Wide Web Conference on World Wide Web, 803-812. International World Wide Web Conferences Steering Committee. (cited on page 17)

Kuss, M. and Rasmussen, C. E., 2005. Assessing approximate inference for binary gaussian process classification. Journal of machine learning research, 6, Oct (2005), 1679-1704. (cited on pages xix, 55, and 56)

Lallouache, M. and Challet, D., 2016. The limits of statistical significance of hawkes processes fitted to financial data. Quantitative Finance, 16, 1 (2016), 1-11. (cited on page 2)

LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P., 1998. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86, 11 (1998), 2278-2324. (cited on page 69)

Lewis, E. and Mohler, G., 2011. A nonparametric em algorithm for multiscale hawkes processes. Journal of Nonparametric Statistics, 1, 1 (2011), 1-20. (cited on pages $2,4,8,17$, and 19)

Lewis, G. and Syrgkanis, V., 2018. Adversarial generalized method of moments. (cited on pages $16,21,22,60,68$, and 121)

Li, S.; Xiao, S.; Zhu, S.; Du, N.; Xie, Y.; and Song, L., 2018. Learning temporal point processes via reinforcement learning. In Advances in Neural Information Processing Systems, vol. 31. Curran Associates, Inc. URL https://proceedings.neurips.cc/paper /2018/file/5d50d22735a7469266aab23fd8aeb536-Paper.pdf. (cited on page 18)

Li, Y.; Hernández-Lobato, J. M.; and Turner, R. E., 2015. Stochastic expectation propagation. In Advances in neural information processing systems, 2323-2331. (cited on pages 3 and 20)

Li, Y. and Turner, R. E., 2016. Rényi divergence variational inference. In Advances in Neural Information Processing Systems, 1073-1081. (cited on pages 19 and 50)

Liao, L.; Chen, Y.; Yang, Z.; Dai, B.; Wang, Z.; and Kolar, M., 2020. Provably efficient neural estimation of structural equation model: An adversarial approach. In Advances in Neural Information Processing Systems 33. Curran Associates, Inc. (cited on page 21)

Liao, L.; Fu, Z.; Yang, Z.; Kolar, M.; and Wang, Z., 2021. Instrumental variable value iteration for causal offline reinforcement learning. arXiv preprint arXiv:2102.09907, (2021). (cited on page 20)

Linderman, S. and Adams, R., 2014. Discovering latent network structure in point process data. In International Conference on Machine Learning, 1413-1421. (cited on pages 3 and 18)

Linderman, S. W. and Adams, R. P., 2015. Scalable bayesian inference for excitatory point process networks. arXiv preprint arXiv:1507.03228, (2015). (cited on pages 3, 18, and 19)

Liu, Q.; Li, L.; Tang, Z.; and Zhou, D., 2018. Breaking the curse of horizon: Infinitehorizon off-policy estimation. In Advances in Neural Information Processing Systems, vol. 31. Curran Associates, Inc. (cited on pages 118 and 119)

Lloyd, C.; Gunter, T.; Osborne, M.; and Roberts, S., 2015. Variational inference for gaussian process modulated poisson processes. In International Conference on Machine Learning, 1814-1822. (cited on pages 4, 23, 24, 35, 36, 37, and 41)

Luenberger, D. G., 1997. Optimization by vector space methods. John Wiley \& Sons. (cited on page 112)

Mallasto, A. and Feragen, A., 2017. Learning from uncertain curves: The 2wasserstein metric for gaussian processes. In Advances in Neural Information Processing Systems 30 (Eds. I. Guyon; U. V. Luxburg; S. Bengio; H. Wallach; R. Fergus; S. Vishwanathan; and R. Garnett), 5660-5670. Curran Associates, Inc. (cited on page 20)

Mann, H. B. and Wald, A., 1943. On stochastic limit and order relationships. The Annals of Mathematical Statistics, 14, 3 (09 1943), 217-226. (cited on page 101)

Matheron, G., 1963. Principles of geostatistics. Economic Geology, 58, 8 (12 1963), 1246-1266. (cited on pages 13 and 50)

Meehan, M. and Penckofer, S., 2014. The role of vitamin D in the aging adult. Journal of aging and gerontology, 2, 2 (2014), 60. (cited on pages xvii, 71, and 126)

Mei, H. and Eisner, J. M., 2017. The neural hawkes process: A neurally selfmodulating multivariate point process. In Advances in Neural Information Processing Systems, 6757-6767. (cited on pages 17 and 18)

Mercer, J., 1909. Functions of positive and negative type, and their connection with the theory of integral equations. Philosophical Transactions of the Royal Society A, 209 (1909), 415-446. URL http://www.jstor.org/sle/91043. (cited on pages 10 and 24)

Minka, T., 2004. Power ep. Dep. Statistics, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep, (2004). (cited on page 19)

Minka, T., 2005. Divergence measures and message passing. Technical report, Microsoft Research. (cited on pages 5, 19, and 50)

Minka, T. P., 2001a. A family of algorithms for approximate Bayesian inference. Ph.D. thesis, Massachusetts Institute of Technology. (cited on page 20)

Minka, T. P., 2001b. The EP energy function and minimization schemes. Technical report, Technical report. (cited on pages 13 and 19)

Minka, T. P., 2001c. Expectation propagation for approximate bayesian inference. In Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence, 362-369. Morgan Kaufmann Publishers Inc. (cited on pages 13, 19, and 50)

Mishra, S.; Rizoiu, M.-A.; and Xie, L., 2016. Feature driven and point process approaches for popularity prediction. In Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, 1069-1078. ACM. (cited on pages 2, 17, and 32)

Møller, J.; Syversveen, A. R.; and Waggepetersen, R. P., 1998. Log gaussian cox processes. Scandinavian journal of statistics, 25, 3 (1998), 451-482. (cited on pages 24 and 50)

Montavon, G.; Müller, K.-R.; and Cuturi, M., 2016. Wasserstein training of restricted boltzmann machines. In Advances in Neural Information Processing Systems, 3718-3726. (cited on page 19)

Muandet, K.; Fukumizu, K.; Sriperumbudur, B.; and Schölkopf, B., 2017. Kernel mean embedding of distributions: A review and beyond. Foundations and Trends in Machine Learning, 10, 1-2 (2017), 1-141. (cited on pages 21 and 61)

Muandet, K.; Jitкrittum, W.; and Kübler, J., 2020a. Kernel conditional moment test via maximum moment restriction. In Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence, vol. 124 of Proceedings of Machine Learning Research, 41-50. PMLR. (cited on pages 22, 60, 61, and 62)

Muandet, K.; Mehrjou, A.; Lee, S. K.; and Raj, A., 2020b. Dual instrumental variable regression. In Advances in Neural Information Processing Systems 33. Curran Associates, Inc. Forthcoming. (cited on pages 16, 20, 21, and 22)

Muth, J. F., 1961. Rational expectations and the theory of price movements. Econometrica: Journal of the Econometric Society, (1961), 315-335. (cited on page 20)

Muzellec, B. and Cuturi, M., 2018. Generalizing point embeddings using the wasserstein space of elliptical distributions. In Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS'18 (Montr\&\#233;al, Canada, 2018), 10258-10269. Curran Associates Inc., USA. (cited on page 20)

Neal, R. M., 1997. Monte carlo implementation of gaussian process models for bayesian regression and classification. arXiv preprint physics/9701026, (1997). (cited on page 50)

Newey, W., 1993. Efficient estimation of models with conditional moment restrictions. In Handbook of Statistics, vol. 11, chap. 16, 419-454. Elsevier. (cited on pages 15, 20, and 60)

Newey, W. K. and McFadden, D., 1994. Chapter 36 large sample estimation and hypothesis testing. In Handbook of Econometrics, vol. 4, 2111 - 2245. Elsevier. (cited on pages $67,97,99,100,101,102,104,105,106$, and 107)

Newey, W. K. and Powell, J. L., 2003. Instrumental variable estimation of nonparametric models. Econometrica, 71, 5 (2003), 1565-1578. (cited on pages 21 and 62)

Newey, W. K. and Smith, R. J., 2004. Higher order properties of gmm and generalized empirical likelihood estimators. Econometrica, 72, 1 (2004), 219-255. (cited on page 61)

Omi, T.; ueda, n.; and Aihara, K., 2019. Fully neural network based model for general temporal point processes. In Advances in Neural Information Processing Systems, vol. 32. Curran Associates, Inc. URL https://proceedings.neurips.cc/paper/2 019/file/39e4973ba3321b80f37d9b55f63ed8b8-Paper.pdf. (cited on pages 17 and 18)

Opper, M. and Archambeau, C., 2009. The variational Gaussian approximation revisited. Neural computation, 21, 3 (2009), 786-792. (cited on pages 50 and 54)

Opper, M. and Winther, O., 2000. Gaussian processes for classification: Mean-field algorithms. Neural computation, 12, 11 (2000), 2655-2684. (cited on pages 3, 13, 19, and 50)

Opper, M. and Winther, O., 2005. Expectation consistent approximate inference. Journal of Machine Learning Research, 6, Dec (2005), 2177-2204. (cited on page 20)

Owen, D. B., 1956. Tables for computing bivariate normal probabilities. Ann. Math. Statist., 27, 4 (12 1956), 1075-1090. doi:10.1214/aoms/1177728074. (cited on page 86)

Ozaki, T., 1979. Maximum likelihood estimation of hawkes' self-exciting point processes. Annals of the Institute of Statistical Mathematics, 31, 1 (1979), 145-155. (cited on page 17)

Pascanu, R.; Mikolov, T.; and Bengio, Y., 2013. On the difficulty of training recurrent neural networks. In Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13 (Atlanta, GA, USA, 2013), III-1310-III-1318. JMLR.org. (cited on page 18)

Peyré, G.; Cuturi, M.; et al., 2019. Computational optimal transport. Foundations and Trends® in Machine Learning, 11, 5-6 (2019), 355-607. (cited on pages 14 and 15)

Quiñonero-Candela, J. and Rasmussen, C. E., 2005. A unifying view of sparse approximate gaussian process regression. Journal of Machine Learning Research, 6, 65 (2005), 1939-1959. URL http://jmlr.org/papers/v6/quinonero-candela05a.html. (cited on page 10)

Rasmussen, C. E. and Williams, C. K. I., 2005. Gaussian Processes for Machine Learning. The MIT Press. ISBN 026218253X. (cited on pages 11, 24, 25,52, 84, 85, 94, and 116)

Rasmussen, J. G., 2013. Bayesian inference for hawkes processes. Methodology and Computing in Applied Probability, 15, 3 (2013), 623-642. (cited on pages 3, 18, and 19)

Ripley, B. D., 1996. Pattern Recognition and Neural Networks. Cambridge University Press. doi:10.1017/CBO9780511812651. (cited on page 55)

Rizoiv, M.-A.; Mishra, S.; Kong, Q.; Carman, M.; and Xie, L., 2017. Sir-Hawkes: on the Relationship Between Epidemic Models and Hawkes Point Processes. arXiv preprint, (nov 2017), 1-16. URL http://arxiv.org/abs/1711.01679. (cited on page 2)

Rizoiv, M.-A.; Mishra, S.; Kong, Q.; Carman, M.; and Xie, L., 2018. Sir-hawkes: Linking epidemic models and hawkes processes to model diffusions in finite populations. In Proceedings of the 27th International Conference on World Wide Web, 419-428. International World Wide Web Conferences Steering Committee. (cited on pages 17,32 , and 46)

Rubin, I., 1972. Regular point processes and their detection. IEEE Transactions on Information Theory, 18, 5 (1972), 547-557. (cited on page 7)

Rubner, Y.; Tomasi, C.; and Guibas, L. J., 2000. The earth mover's distance as a metric for image retrieval. International journal of computer vision, 40, 2 (2000), 99-121. (cited on page 50)

Rue, H.; Martino, S.; and Chopin, N., 2009. Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 71, 2 (2009), 319-392. doi:https://doi.org/10.1111/j.1467-9868.2008.00700.x. URL https: //rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9868.2008.00700.x. (cited on pages 33 and 34)

Rue, H.; Riebler, A.; Sørbye, S. H.; Illian, J. B.; Simpson, D. P.; and Lindgren, F. K., 2017. Bayesian computing with inla: A review. Annual Review of Statistics and Its Application, 4, 1 (2017), 395-421. doi:10.1146/annurev-statistics-060116-054045. URL https://doi.org/10.1146/annurev-statistics-060116-054045. (cited on page 33)

Schölкopf, B.; Herbrich, R.; and Smola, A. J., 2001a. A generalized representer theorem. In Computational Learning Theory, 416-426. Springer Berlin Heidelberg, Berlin, Heidelberg. (cited on page 22)

Schölкopf, B.; Herbrich, R.; and Smola, A. J., 2001b. A generalized representer theorem. In International conference on computational learning theory, 416-426. Springer. (cited on page 64)

Schölкорғ, B. and Smola, A., 2002. Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond. MIT Press, Cambridge, MA, USA. (cited on page 61)

Seeger, M., 2005. Expectation propagation for exponential families. Technical report, Department of EECS, University of California at Berkeley. (cited on pages 13, 19, and 53)

Serfling, R., 1980. Approximation theorems of mathematical statistics. John Wiley \& Sons. (cited on pages 60, 63, 103, and 106)

Shafieezadeh-Abadeh, S.; Nguyen, V. A.; Kuhn, D.; and Esfahani, P. M., 2018. Wasserstein distributionally robust kalman filtering. In Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS'18 (Montr\&\#233;al, Canada, 2018), 8483-8492. Curran Associates Inc., USA. (cited on page 20)

Shchur, O.; Biloš, M.; and Günnemann, S., 2020a. Intensity-free learning of temporal point processes. International Conference on Learning Representations (ICLR), (2020). (cited on page 18)

Shchur, O.; Gao, N.; Biloš, M.; and Günnemann, S., 2020b. Fast and flexible temporal point processes with triangular maps. In Advances in Neural Information Processing Systems, vol. 33, 73-84. Curran Associates, Inc. URL https://proceedings. neurips.cc/paper/2020/file/00ac8ed3b4327bdd4ebbebcb2ba10a00-Paper.pdf. (cited on page 18)

Shirai, T. and Takahashi, Y., 2003. Random point fields associated with certain fredholm determinants ii: Fermion shifts and their ergodic and gibbs properties. The Annals of Probability, 31, 3 (2003), 1533-1564. (cited on page 24)

Simma, A. and Jordan, M. I., 2010. Modeling events with cascades of poisson processes. In Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, UAI'10 (Catalina Island, CA, 2010), 546-555. AUAI Press, Arlington, Virginia, United States. URL http://dl.acm.org/citation.cfm?id=3023549.3023614. (cited on page 1)

Simon-Gabriel, C.-J. and Schölкopf, B., 2018. Kernel distribution embeddings: Universal kernels, characteristic kernels and kernel metrics on distributions. Journal of Machine Learning Research, 19, 44 (2018), 1-29. (cited on pages 62 and 97)

Singh, R.; Sahani, M.; and Gretton, A., 2019. Kernel instrumental variable regression. In Advances in Neural Information Processing Systems 32, 4593-4605. Curran Associates, Inc. (cited on pages 16, 21, 22, 68, and 120)

Sjolander, A. and Martinussen, T., 2019. Instrumental variable estimation with the R package ivtools. Epidemiologic Methods, 8, 1 (2019). (cited on pages xvii, 71, 124, and 126)

Snelson, E. and Ghahramani, Z., 2006. Sparse gaussian processes using pseudoinputs. In Advances in neural information processing systems, 1257-1264. (cited on page 41)

Snelson, E.; Ghahramani, Z.; and Rasmussen, C. E., 2004. Warped gaussian processes. In Advances in Neural Information Processing Systems 16 (Eds. S. Thrun; L. K. Saul; and B. Scholкорf), 337-344. MIT Press. (cited on page 13)

Song, L.; Fukumizu, K.; and Gretton, A., 2013. Kernel embeddings of conditional distributions: A unified kernel framework for nonparametric inference in graphical models. IEEE Signal Processing Magazine, 30, 4 (2013), 98-111. (cited on page 21)

Song, L.; Huang, J.; Smola, A.; and Fukumizu, K., 2009. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In Proceedings of the 26th International Conference on Machine Learning (ICML). (cited on page 21)

Sriperumbudur, B. K.; Fukumizu, K.; and Lanckriet, G. R. G., 2011. Universality, characteristic kernels and rkhs embedding of measures. Journal of Machine Learning Research, 12 (Jul. 2011), 2389-2410. (cited on page 97)

Steinwart, I., 2002. On the influence of the kernel on the consistency of support vector machines. Journal of Machine Learning Research, 2 (Mar. 2002), 67-93. (cited on page 97)

Steinwart, I. and Christmann, A., 2008. Support vector machines. Springer Science \& Business Media. (cited on pages 97, 98, and 108)

Takatsu, A., 2011. Wasserstein geometry of gaussian measures. Osaka J. Math., 48, 4 (12 2011), 1005-1026. (cited on page 20)

Zhang, R.; Imaizumi, M.; Schölкорғ, B.; and Muandet, K., 2021a. Maximum moment restriction for instrumental variable regression. arXiv preprint, (2021). Submitted to NeurIPS 2021. (cited on pages ix and 3)

Zhang, R.; Muandet, K.; Schölкорf, B.; and Imaizumi, M., 2021b. Instrument space selection for kernel maximum moment restriction. arXiv preprint, (2021). Submitted to NeurIPS 2021. (cited on page ix)

Zhang, R.; Walder, C.; Bonilla, E. V.; Rizoiu, M.-A.; and Xie, L., 2020a. Quantile propagation for wasserstein-approximate gaussian processes. In Advances in Neural Information Processing Systems (NeurIPS 2020), vol. 33, 21566-21578. Curran Associates, Inc. (cited on pages ix and 3)

Zhang, R.; Walder, C.; and Rizoiu, M.-A., 2020b. Variational inference for sparse gaussian process modulated hawkes proces. In The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020). New York, U.S.A. (cited on pages ix, 3 , and 94 )

Zhang, R.; Walder, C.; Rizoiv, M.-A.; and Xie, L., 2019. Efficient non-parametric bayesian hawkes processes. In International Joint Conference on Artificial Intelligence (IJCAI 2019). Macao, China. (cited on pages ix, $3,42,43$, and 50)

Titsias, M. K., 2009. Variational model selection for sparse gaussian process regression. Technical report, Technical report, School of Computer Science, University of Manchester. (cited on pages 10 and 11)

Tolstikhin, I.; Bousquet, O.; Gelly, S.; and Schoelkopf, B., 2017. Wasserstein auto-encoders. arXiv preprint arXiv:1711.01558, (2017). (cited on page 51)

Tsүbaкov, A., 2008. Introduction to Nonparametric Estimation. Springer Publishing Company, Incorporated, 1st edn. (cited on page 21)

Uehara, M.; Huang, J.; and Jiang, N., 2020. Minimax weight and q-function learning for off-policy evaluation. In International Conference on Machine Learning, 9659-9668. PMLR. (cited on page 118)

Van der Vaart, A., 2000. Asymptotic Statistics. Cambridge University Press. (cited on pages 64 and 101)

Van der Vaart, A. and Wellner, J., 1996. Weak convergence and empirical processes. (cited on page 110)

Vapnik, V. N., 1998. Statistical Learning Theory. Wiley-Interscience. (cited on page 63)
Vehtari, A.; Mononen, T.; Tolvanen, V.; Sivula, T.; and Winther, O., 2016. Bayesian leave-one-out cross-validation approximations for Gaussian latent variable models. The Journal of Machine Learning Research, 17, 1 (2016), 3581-3618. (cited on page 66)

Villani, C., 2008. Optimal transport: old and new, vol. 338. Springer Science \& Business Media. (cited on page 50)

Walder, C. J. and Bishor, A. N., 2017. Fast bayesian intensity estimation for the permanental process. In International Conference on Machine Learning, 3579-3588. (cited on pages $4,18,23,24,26,27,56,86,94$, and 115)

Williams, C. K. I. and Barber, D., 1998. Bayesian classification with gaussian processes. IEEE Transactions on Pattern Analysis and Machine Intelligence, 20, 12 (Dec 1998), 1342-1351. doi:10.1109/34.735807. (cited on pages 13 and 50)

Williams, C. K. I. and Seeger, M., 2001. Using the nyström method to speed up kernel machines. In Advances in Neural Information Processing Systems 13 (Eds. T. K. Leen; T. G. Dietterich; and V. Tresp), 682-688. Mit Press. (cited on pages 43 and 65)

Wolfram, R. I., 2019. Mathematica, Version 12.0. Champaign, IL, 2019. (cited on pages 87 and 88)

Wong, R. K. and Chan, K. C. G., 2018. Kernel-based covariate functional balancing for observational studies. Biometrika, 105, 1 (2018), 199-213. (cited on page 118)

Wu, F. and Huberman, B. A., 2007. Novelty and collective attention. PNAS, 104, 45 (nov 2007), 17599-601. doi:10.1073/pnas.0704916104. URL http://www.pnas.org/con tent/104/45/17599.abstract. (cited on page 2)

Xiao, S.; Farajtabar, M.; Ye, X.; Yan, J.; Song, L.; and Zha, H., 2017a. Wasserstein learning of deep generative point process models. In Advances in Neural Information Processing Systems, 3247-3257. (cited on pages 17 and 18)

Xiao, S.; Yan, J.; Yang, X.; Zha, H.; and Chu, S., 2017b. Modeling the intensity function of point process via recurrent neural networks. In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 31. (cited on pages 17 and 18)

Xu, H.; Farajtabar, M.; and Zha, H., 2016. Learning granger causality for hawkes processes. In Proceedings of The 33rd International Conference on Machine Learning, vol. 48 of Proceedings of Machine Learning Research, 1717-1726. PMLR, New York, New York, USA. (cited on pages 17 and 19)

Xu, H.; Luo, D.; Chen, X.; and Carin, L., 2018. Benefits from superposed hawkes processes. International Conference on Artificial Intelligence and Statistics (AISTATS), (2018), 623-631. (cited on page 30)

Xu, M.; Lakshminarayanan, B.; Teh, Y. W.; Zhu, J.; and Zhang, B., 2014. Distributed bayesian posterior sampling via moment sharing. In Advances in Neural Information Processing Systems. (cited on page 20)

Zhang, Q.; Lipani, A.; Kirnap, O.; and Yilmaz, E., 2020. Self-attentive Hawkes process. In Proceedings of the 37th International Conference on Machine Learning, vol. 119 of Proceedings of Machine Learning Research, 11183-11193. PMLR. URL http://proceedings.mlr.press/v119/zhang20q.html. (cited on pages 17 and 18)

Zhao, Q.; Erdogdu, M. A.; He, H. Y.; Rajaraman, A.; and Leskovec, J., 2015. Seismic: A self-exciting point process model for predicting tweet popularity. In Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1513-1522. ACM. (cited on pages 1, 32, and 46)

Zhou, D.-X., 2002. The covering number in learning theory. Journal of Complexity, 18, 3 (2002), 739-767. (cited on page 108)

Zhou, K.; Zha, H.; and Song, L., 2013. Learning triggering kernels for multidimensional hawkes processes. In International Conference on Machine Learning, 1301-1309. (cited on pages $2,4,17,19,29,43$, and 45)

Zou, G., 2004. A modified poisson regression approach to prospective studies with binary data. American journal of epidemiology, 159, 7 (2004), 702-706. (cited on page 13)

Zuo, S.; Jiang, H.; Li, Z.; Zhao, T.; and Zha, H., 2020. Transformer hawkes process. In International Conference on Machine Learning, 11692-11702. PMLR. (cited on pages 17 and 18)


[^0]:    ${ }^{1} \operatorname{Gamma}\left(\mu \mid k_{0}, c_{0}\right)=\frac{1}{\Gamma\left(k_{0}\right) c_{0}^{k}} \mu^{k_{0}-1} e^{-x / c_{0}}$

[^1]:    ${ }^{2}$ The CPU we use is Intel(R) Core(TM) i7-7700 CPU @ 3.60 GHz and the language is Python 3.6.5.

[^2]:    ${ }^{1}$ Also, IV assumptions imply $\varepsilon$ is independent of $Z$, i.e., $\mathbb{E}[(Y-f(X)) h(Z)]=0$ for all measurable $h$.

[^3]:    ${ }^{1} \operatorname{Gamma}(x \mid k, c)=\frac{1}{\Gamma(k) c^{k}} x^{k-1} e^{-x / c}$.

[^4]:    ${ }^{1}$ Throughout the paper, we denote by $N\left(\mu, \sigma^{2}\right)$ a Gaussian distribution with the mean $\mu$ and variance $\sigma^{2}$.

