
Importance Weighting Approach in Kernel Bayes' Rule

Liyuan Xu¹ Yutian Chen² Arnaud Doucet² Arthur Gretton¹

Abstract

We study a nonparametric approach to Bayesian computation via feature means, where the expectation of prior features is updated to yield expected kernel posterior features, based on regression from learned neural net or kernel features of the observations. All quantities involved in the Bayesian update are learned from observed data, making the method entirely model-free. The resulting algorithm is a novel instance of a kernel Bayes' rule (KBR). Our approach is based on importance weighting, which results in superior numerical stability to the existing approach to KBR, which requires operator inversion. We show the convergence of the estimator using a novel consistency analysis on the importance weighting estimator in the infinity norm. We evaluate our KBR on challenging synthetic benchmarks, including a filtering problem with a state-space model involving high dimensional image observations. The proposed method yields uniformly better empirical performance than the existing KBR, and competitive performance with other competing methods.

1. Introduction

Many machine learning applications are reduced to the problem of inferring latent variables in probabilistic models. This is achieved using Bayes' rule, where a prior distribution over the latent variables is updated to obtain the posterior distribution, based on a likelihood function. Probabilities involved in the Bayes' update have generally been expressed as probability density functions. For many interesting problems, the exact posterior density is intractable; in the event that the likelihood is known, we may use approximate inference techniques, including *Markov Chain Monte Carlo (MCMC)* and *Variational inference (VI)*. When the

likelihood function is unknown, however, and only samples drawn from it can be obtained, these methods do not apply.

We propose to use a *kernel mean embedding* representation of our probability distributions as the key quantity in our Bayesian updates (Smola et al., 2007), which will enable nonparametric inference without likelihood functions. Kernel mean embeddings characterize probability distributions as expectations of features in a reproducing kernel Hilbert space (RKHS). When the space is *characteristic*, these embeddings are injective (Fukumizu et al., 2009; Sriperumbudur et al., 2010). Expectations of RKHS functions may be expressed as dot products with the corresponding mean embeddings. Kernel mean embeddings have been employed extensively in nonparametric hypothesis testing (Gretton et al., 2012; Chwialkowski et al., 2016), however they have also been used in supervised learning for distribution-valued inputs (Muandet et al., 2012; Szabo et al., 2016), and in various other inference settings (Song et al., 2009; Grunewalder et al., 2012; Boots et al., 2013; Jitkrittum et al., 2015; Singh et al., 2019; Muandet et al., 2021).

We will focus here on the *Kernel Bayes' Rule (KBR)*, where a prior mean embedding is taken as input and is updated to return the mean embedding of the posterior. The goal is to express all quantities involved in the KBR updates as RKHS functions learned from observed data: parametric models should not be required at any stage in the computation. We will address in particular the application of kernel Bayes' rule to filtering, where a latent state evolves stochastically over time according to Markovian dynamics, with noisy, high dimensional observations providing information as to the current state: the task is to construct a posterior distribution over the current state from the sequence of noisy observations up to the present. In the event that a parametric model is known for the latent dynamics, then filtering can be achieved either by sampling from the model (Kanagawa et al., 2016), or through closed-form updates (Nishiyama et al., 2020), however, we propose here to make no modeling assumptions concerning the latent dynamics, beyond observing them in a training stage. An alternative to kernel Bayes updates in the filtering setting was proposed by Jasra et al. (2012) but, for high-dimensional observations, it requires introducing summary statistics which have a high impact on performance and are difficult to select.

¹Gatsby Unit ²DeepMind. Correspondence to: Liyuan Xu <liyuan.jo.19@ucl.ac.uk>.

The first formulation of Kernel Bayes' Rule, due to Fukumizu et al. (2013), yields the desired model-free, sample-derived posterior embedding estimates, and a filtering procedure that is likewise nonparametric and model-free. At a high level, the resulting posterior is obtained via a two-stage least squares procedure, where the output of the first regression stage is used as input to the second stage. Unfortunately, the original KBR formulation requires an unconventional form of regularization in the second stage, which adversely impacts the attainable convergence rates (see Section 3.2), and decreases performance in practice.

In the present work, we introduce *importance-weighted KBR (IW-KBR)*, a novel design for a KBR which does not require the problematic second-stage regularization. The core idea of our approach is to use importance weighting, rather than two-stage regression, to achieve the required Bayesian update. We provide a convergence analysis, requiring as an intermediate step a novel analysis on our importance weighting estimator in infinity norm, which may be of independence interest. Our IW-KBR improves on the convergence rate of the original KBR under certain conditions. As an algorithmic contribution, we introduce adaptive neural network features into IW-KBR, which is essential when the observations are images or other complex objects. In experiments, IW-KBR outperforms the original KBR across all benchmarks considered, including filtering problems with state-space models involving high dimensional image observations.

The paper is structured as follows. In Section 2, we introduce the basic concepts of kernel methods and review the original KBR approach. We will also introduce density ratio estimators, which are used in estimating the importance weights. Then, in Section 3 we introduce the proposed KBR approach using kernel and neural net features and provide convergence guarantees when kernel features are employed. We describe the *Kernel Bayes Filter (KBF)* in Section 4, which applies KBR to the filtering problem in a state-space model. We demonstrate the empirical performance in Section 6, covering three settings: non-parametric Bayesian inference, low-dimensional synthetic settings introduced in Fukumizu et al. (2013), and challenging KBF settings where the observations consist of high-dimensional images.

2. Background and preliminaries

In this section, we introduce kernel mean embeddings, which represent probability distributions by expected RKHS features. We then review Kernel Bayes' Rule (KBR) (Fukumizu et al., 2013), which aims to learn a mean posterior embedding according to Bayes' rule.

Kernel mean embeddings: Let (X, Z) be random variables on $\mathcal{X} \times \mathcal{Z}$ with distribution P and density function $p(x, z)$,

and $k_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $k_{\mathcal{Z}} : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ be measurable positive definite kernels corresponding to the scalar-valued RKHSs $\mathcal{H}_{\mathcal{X}}$ and $\mathcal{H}_{\mathcal{Z}}$, respectively. We denote the feature maps as $\psi(x) = k_{\mathcal{X}}(x, \cdot)$ and $\phi(z) = k_{\mathcal{Z}}(z, \cdot)$, and RKHS norm as $\|\cdot\|$.

The kernel mean embedding $m_{P(X)}$ of a marginal distribution $P(X)$ is defined as

$$m_{P(X)} = \mathbb{E}_P[\psi(X)] \in \mathcal{H}_{\mathcal{X}},$$

and always exists for bounded kernels. From the reproducing property, we have $\langle f, m_{P(X)} \rangle = \mathbb{E}_P[f(X)]$ for all $f \in \mathcal{H}_{\mathcal{X}}$, which is useful to estimate the expectations of functions. Furthermore, it is known that the embedding uniquely defines the probability distribution (i.e. $m_{P(X)} = m_{Q(X)}$ implies $P = Q$) if kernel $k_{\mathcal{X}}$ is *characteristic*: for instance, the Gaussian kernel has this property (Fukumizu et al., 2009; Sriperumbudur et al., 2010). In addition to kernel means, we will require kernel covariance operators,

$$C_{XX}^{(P)} = \mathbb{E}_P[\psi(X) \otimes \psi(X)], \quad C_{XZ}^{(P)} = \mathbb{E}_P[\psi(X) \otimes \phi(Z)], \\ C_{ZX}^{(P)} = \left(C_{XZ}^{(P)}\right)^*, \quad C_{ZZ}^{(P)} = \mathbb{E}_P[\phi(Z) \otimes \phi(Z)],$$

where \otimes is the tensor product such that $[a \otimes b]c = a \langle b, c \rangle$, and $*$ denotes the adjoint of the operator. Covariance operators generalize finite-dimensional covariance matrices to the case of infinite kernel feature spaces, and always exist for bounded kernels.

The *kernel conditional mean embedding* (Song et al., 2009) is the extension of the kernel mean embedding to conditional probability distributions, and is defined as

$$m_{P(Z|X)}(x) = \mathbb{E}_P[\phi(Z)|X = x] \in \mathcal{H}_{\mathcal{Z}}.$$

Under the regularity condition $\mathbb{E}_P[g(Z)|X = x] \in \mathcal{H}_{\mathcal{Z}}$ for all $g \in \mathcal{H}_{\mathcal{Z}}$, there exists an operator $E_P \in \mathcal{H}_{\mathcal{X}} \otimes \mathcal{H}_{\mathcal{Z}}$ such that $m_{P(Z|X)}(x) = (E_P)^* \psi(x)$ (Song et al., 2009; Grunewalder et al., 2012; Singh et al., 2019). The *conditional operator* E_P can be expressed as the minimizer of the surrogate regression loss $\ell^{(P)}$, defined as

$$\ell^{(P)}(E) = \mathbb{E}_P[\|\phi(Z) - E^* \psi(X)\|^2].$$

This minimization can be solved analytically, and the closed-form solution is given as

$$E_P = (C_{XX}^{(P)})^{-1} C_{XZ}^{(P)}.$$

An empirical estimate of the conditional mean embedding is straightforward. Given i.i.d samples $\{(x_i, z_i)\}_{i=1}^n \sim P$, we minimize $\hat{\ell}^{(P, \lambda)}$, defined as

$$\hat{\ell}^{(P, \lambda)}(E) = \frac{1}{n} \sum_{i=1}^n \|\phi(z_i) - E^* \psi(x_i)\|^2 + \lambda \|E\|^2,$$

which is the sample estimate of $\ell^{(P)}$ with added Tikhonov regularization, and the norm $\|E\|$ is the Hilbert-Schmidt

norm. The solution of this minimization problem is

$$\hat{E}_{P,\lambda} = (\hat{C}_{XX}^{(P)} + \lambda I)^{-1} \hat{C}_{XZ}^{(P)}.$$

Here, we denote $\hat{C}_{ZX}^{(P)}, \hat{C}_{XX}^{(P)}, \hat{C}_{XZ}^{(P)}, \hat{C}_{ZZ}^{(P)}$ as the empirical estimates of the covariance operators $C_{ZX}^{(P)}, C_{XX}^{(P)}, C_{XZ}^{(P)}, C_{ZZ}^{(P)}$, respectively; e.g. $\hat{C}_{ZX}^{(P)} = \frac{1}{n} \sum_{i=1}^n \phi(z_i) \otimes \psi(x_i)$.

Kernel Bayes' Rule: Based on conditional mean embedding, Fukumizu et al. (2013) proposed a method to realize Bayes' rule in kernel mean embedding representation. In Bayes' rule, we aim to update the prior distribution Π on latent variables Z , with the density function $\pi(z)$, to the posterior $Q(Z|X = \tilde{x})$, where the joint density $q(z, x)$ of distribution Q is given by

$$q(z, x) = \pi(z)g(x|z).$$

Here, $g(x|z)$ is the *likelihood* function and $\tilde{x} \in \mathcal{X}$ is the *conditioning point*. The density of $Q(Z|X = \tilde{x})$ is often intractable, since it involves computing the integral $\int q(z, x) dz$. Instead, KBR aims to update the embedding of Π , denoted as m_Π , to the embedding of the posterior $m_{Q(Z|X)}(\tilde{x})$. Fukumizu et al. (2013) show that such an update does not require the closed-form expression of likelihood function $g(x|z)$. Rather, KBR *learns* the relations between latent and observable variables from the data $\{(x_i, z_i)\}_{i=1}^n \sim P$, where the density of data distribution P shares the same likelihood $p(z, x) = p(z)g(x|z)$. We require P to share the likelihood of Q , otherwise the relation between observations and latent cannot be learned. We remark that KBR also applies to *Approximate Bayesian Computation (ABC)* (Tavaré et al., 1997; Marjoram et al., 2003), in which the likelihood function is intractable, but we can simulate from it.

If $\pi(z) = p(z)$, we could simply estimate $m_{Q(Z|X)}(\tilde{x})$ by the conditional mean embedding learned from $\{(x_i, z_i)\}_{i=1}^n$, but we will generally require $\pi(z) \neq p(z)$, since the prior can be the result of another inference step (notably for filtering: see Section 4). We present the solution of Fukumizu et al. (2013) for this case. Let $C_{ZX}^{(Q)}, C_{XX}^{(Q)}, C_{XZ}^{(Q)}, C_{ZZ}^{(Q)}$ be the covariance operators on distribution Q , defined similarly to the covariance operators on P . Similarly, the conditional operator $E_Q : \mathcal{H}_X \otimes \mathcal{H}_Z$ satisfying $m_{Q(Z|X)}(x) = E_Q^* \psi(x)$ is a minimizer of the loss

$$\ell^{(Q)}(E) = \mathbb{E}_Q [\|\phi(Z) - E^* \psi(X)\|^2].$$

Unlike in the conditional mean embedding case, however, we cannot directly minimize this loss, since data $\{x_i, z_i\}$ are not sampled from Q . Instead, Fukumizu et al. uses the analytical form of E_Q :

$$E_Q = C_{ZX}^{(Q)}(C_{XX}^{(Q)})^{-1},$$

and replaces each operator with vector-valued kernel ridge

regression estimates,

$$\begin{aligned} \hat{C}_{XZ}^{(Q,\text{org})} &= \sum_{i=1}^n \gamma_i \psi(x_i) \otimes \phi(z_i), \\ \hat{C}_{XX}^{(Q,\text{org})} &= \sum_{i=1}^n \gamma_i \psi(x_i) \otimes \psi(x_i), \end{aligned}$$

where each weight γ_i is given as

$$\gamma_i = \left\langle \phi(z_i), \left(\hat{C}_{ZZ}^{(P)} + \eta I \right)^{-1} \hat{m}_\Pi \right\rangle. \quad (1)$$

Here, η is another Tikhonov regularization parameter. Although these estimators are consistent, $\hat{C}_{XX}^{(Q,\text{org})}$ is not necessarily positive semi-definite since weight γ_i can be negative. This causes instabilities when inverting the operator. Fukumizu et al. mitigate this by applying another type of Tikhonov regularization, yielding an alternative estimate of E_Q ,

$$\hat{E}_{Q,\lambda}^{(\text{org})} = \hat{C}_{XX}^{(Q,\text{org})} \left(\left(\hat{C}_{XX}^{(Q,\text{org})} \right)^2 + \lambda I \right)^{-1} \hat{C}_{XZ}^{(Q,\text{org})}. \quad (2)$$

Density Ratio Estimation: The core idea of our proposed approach is to use importance sampling to estimate $\ell^{(Q)}$. To obtain the weights, we need to estimate the density ratio

$$r_0(z) = \pi(z)/p(z).$$

We may use any density ratio estimator, as long as it is computable from data $\{z_i\}_{i=1}^n$ and the prior embedding estimator \hat{m}_Π . We focus here on the *Kernel-Based unconstrained Least-Squares Importance Fitting (KuLSIF)* estimator (Kanamori et al., 2012), which is obtained by minimizing

$$\tilde{r}_\eta = \arg \min_{r \in \mathcal{H}_Z} \frac{1}{2} \left\langle r, \hat{C}_{ZZ}^{(P)} r \right\rangle - \langle r, \hat{m}_\Pi \rangle + \frac{\eta}{2} \|r\|^2.$$

The estimator $\hat{r}(z)$ is obtained by truncating the solution \tilde{r} at zero: $\hat{r}(z) = \max(0, \tilde{r}_\eta(z))$. Interestingly, the KuLSIF estimator at data point z_i can be written

$$\hat{r}(z_i) = \max(0, \gamma_i),$$

where γ_i is the weight used in (1). Kanamori et al. (2012) developed a convergence analysis of this KuLSIF estimator in ℓ_2 -norm based on the bracketing entropy (Cucker & Smale, 2001). This analysis, however, is insufficient for the KBR case, and we will establish stronger convergence results under different assumptions in Section 3.2.

3. Importance-Weighted KBR

In this section, we introduce our *importance-weighted KBR (IW-KBR)* approach and provide a convergence analysis. We also propose a method to learn adaptive features using neural networks so that the model can learn complex posterior distributions.

3.1. Importance Weighted KBR

Our proposed method minimizes the loss $\ell^{(Q)}$, which is estimated by the importance sampling. Using density ratio $r_0(z) = \pi(z)/p(z)$, the loss $\ell^{(Q)}$ can be rewritten as

$$\ell^{(Q)}(E) = \mathbb{E}_P [r_0(Z) \|\phi(Z) - E^* \psi(X)\|^2].$$

Hence, we can construct the empirical loss with added Tikhonov regularization,

$$\hat{\ell}^{(Q,\lambda)}(E) = \frac{1}{n} \sum_{i=1}^n \hat{r}_i \|\phi(z_i) - E^* \psi(x_i)\|^2 + \lambda \|E\|^2, \quad (3)$$

where $\hat{r}_i = \hat{r}(z_i)$ is a non-negative estimator of density ratio $r_0(z_i)$. Again, the minimizer of $\hat{\ell}^{(Q,\lambda)}(E)$ can be obtained analytically as

$$\hat{E}_{Q,\lambda} = \left(\hat{C}_{XX}^{(Q)} + \lambda I \right)^{-1} \hat{C}_{XZ}^{(Q)}, \quad (4)$$

where

$$\hat{C}_{XX}^{(Q)} = \sum_{i=1}^n \frac{\hat{r}_i}{n} \psi(x_i) \otimes \phi(z_i), \quad \hat{C}_{XZ}^{(Q)} = \sum_{i=1}^n \frac{\hat{r}_i}{n} \psi(x_i) \otimes \psi(x_i).$$

Note that $\hat{C}_{XX}^{(Q)}$ is always positive semi-definite since \hat{r}_i is non-negative by definition. Using the KuLSIF estimator, this is the truncated weight $\hat{r}_i = \max(0, \gamma_i)$ described previously.

Given estimated conditional operator $\hat{E}_{Q,\lambda}$ in (4), we can estimate the conditional embedding $m_{Q(Z|X)}(x)$ as $\hat{m}_{Q(Z|X)}(x) = (\hat{E}_{Q,\lambda})^* \psi(x)$, as shown in Algorithm 1. As illustrated, the posterior mean embedding $\hat{m}_{Q(Z|X)}$ is represented by the weighted sum over the same RKHS features $\phi(z) = k_{\mathcal{Z}}(z, \cdot)$ used in the density ratio estimator \hat{r} . We remark that this need not be the case, however: the weights w could be used to obtain a posterior mean embedding over a *different* feature space to that used in computing \hat{r} , simply by substituting the desired $\tilde{\phi}(z_i)$ for $\phi(z_i)$ in the sum from the third step. For example, we could use a Gaussian kernel to compute the density ratio \hat{r} , and then a linear kernel $\tilde{\phi}(z) = z$ to estimate the posterior mean $\hat{\mathbb{E}}_Q[Z|X] = \sum_i w_i z_i$.

The computational complexity of Algorithm 1 is $O(n^3)$, which is the same complexity as the ordinary kernel ridge regression. This can be accelerated by using Random Fourier Features (Rahimi & Recht, 2008) or the Nyström Method (Williams & Seeger, 2001).

3.2. Convergence Analysis

In this section, we analyze $\hat{E}_{Q,\lambda}$ in (4). First, we state a few regularity assumptions.

Assumption 3.1. The kernels $k_{\mathcal{X}}$ and $k_{\mathcal{Z}}$ are continuous and bounded: $\|\phi(z)\| \leq \kappa_1 < \infty$, $\|\psi(x)\| \leq \kappa_2 < \infty$ hold almost surely.

Algorithm 1 Importance Weighted Kernel Bayes Rule

Input: Samples $\{(x_i, z_i)\}_{i=1}^n \sim P$, Estimated prior embedding \hat{m}_{Π} , regularization parameters (η, λ) , Conditioning point \tilde{x} .

1: Compute Gram matrices $G_X, G_Z \in \mathbb{R}^{n \times n}$.

$$(G_X)_{ij} = k_{\mathcal{X}}(x_i, x_j), (G_Z)_{ij} = k_{\mathcal{Z}}(z_i, z_j).$$

2: Compute KuLSIF $\hat{r} = (\hat{r}_1, \dots, \hat{r}_n) \in \mathbb{R}^n$ as

$$\hat{r} = \max \left(0, n (G_Z + n\eta I)^{-1} \mathbf{g}_{\Pi} \right),$$

where max operates element-wisely and

$$(\mathbf{g}_{\Pi})_i = \langle \hat{m}_{\Pi}, \phi(z_i) \rangle.$$

3: Estimate $\hat{m}_{Q(Z|X)}(x) = \sum_{i=1}^n w_i \phi(z_i)$, with weight $w = (w_1, \dots, w_n)$ given as

$$w = \sqrt{D} \left(\sqrt{D} G_X \sqrt{D} + n\lambda I \right)^{-1} \sqrt{D} \mathbf{g}_{\tilde{x}},$$

where $D = \text{diag}(\hat{r}) \in \mathbb{R}^{n \times n}$ and $(\mathbf{g}_{\tilde{x}})_i = k_{\mathcal{X}}(x_i, \tilde{x})$.

Assumption 3.2. We have $r_0 \in \mathcal{H}_{\mathcal{Z}}$ and $\exists g_1 \in \mathcal{H}_{\mathcal{Z}}$ such that $r_0 = (C_{ZZ}^{(P)})^{\beta_1} g_1$ and $\|g_1\| \leq \zeta_1$ for $\beta_1 \in (0, 1/2]$, $\zeta_1 < \infty$ given.

Assumption 3.3. $\exists \Gamma \in \mathcal{H}_{\mathcal{X}} \otimes \mathcal{H}_{\mathcal{Z}}$ such that $E_Q = (C_{XX}^{(Q)})^{\beta_2} \Gamma$ and $\|\Gamma\| \leq \zeta_2$ for $\beta_2 \in (0, 1/2]$, $\zeta_2 < \infty$ given.

Assumption 3.1 is standard for mean embeddings, and many widely used kernel functions satisfy it, including the Gaussian kernel and the Matérn kernel. Assumptions 3.2 and 3.3 assure the smoothness of the density ratio r_0 and conditional operator E_Q , respectively. See (Smale & Zhou, 2007; Caponnetto & Vito, 2007) for a discussion of the first assumption, and (Singh et al., 2019, Hypothesis 5) for the second. We will further discuss the implication of Assumption 3.2 using the ratio of two Gaussian distributions in Appendix A. Note that Assumption 3.2, also used in Fukumizu et al. (2013), should be treated with care. For example, $r_0 \in \mathcal{H}_{\mathcal{Z}}$ imposes $\sup_{z \in \mathcal{Z}} r_0(z) \leq \|r_0\| \sup_{z \in \mathcal{Z}} \|\phi(z)\| < \infty$, which is violated when we consider the ratio of two Gaussian distributions with the different means and the same variance.

We now show that the KuLSIF estimator converges in infinity norm $\|\cdot\|_{\infty}$.

Theorem 3.4. Suppose Assumptions 3.1 and 3.2. Given data $\{x_i, z_i\} \sim P$ and the estimated prior embedding \hat{m}_{Π} such that $\|\hat{m}_{\Pi} - \mathbb{E}_{\Pi}[\phi(X)]\| \leq O_p(n^{-\alpha_1})$ for $\alpha_1 \in (0, 1/2]$, by setting $\eta = O(n^{-\frac{\alpha_1 \beta_1}{\beta_1 + 1}})$, we have

$$\|r_0 - \hat{r}\|_{\infty} \leq O_p \left(n^{-\frac{\alpha_1 \beta_1}{\beta_1 + 1}} \right).$$

The proof is given in Appendix B.1. This result differs from the analysis in Kanamori et al. (2012), which establishes

convergence of the KuLSIF estimator in ℓ_2 -norm, based on the bracketing entropy (Cucker & Smale, 2001). Our result cannot be directly compared with theirs, due to the differences in the assumptions made and the norms used in measuring convergence (in particular, we require convergence in infinity norm). Establishing a relation between the two results represents an interesting research direction, though it is out of the scope of this paper.

Theorem 3.4 can then be used to obtain the convergence rate of covariance operators.

Corollary 3.5. *Given the same conditions in Theorem 3.4, by setting $\eta = O(n^{-\frac{\alpha_1}{\beta_1+1}})$, we have*

$$\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \leq O_p\left(n^{-\frac{\alpha_1\beta_1}{\beta_1+1}}\right).$$

The cross covariance operator $\hat{C}_{XZ}^{(Q)}$ also converges at the same rate. This rate is slower than the original KBR estimator, however, which satisfies

$$\|\hat{C}_{XX}^{(Q,\text{org})} - C_{XX}^{(Q)}\| \leq O_p\left(n^{-\frac{\alpha_1(2\beta_1+1)}{2\beta_1+2}}\right)$$

(Fukumizu et al., 2013). This is inevitable since the original KBR estimator uses the optimal weights γ_i to estimate $C_{XX}^{(Q)}$, while our estimator uses truncated weights $\hat{r}_i = \max(0, \gamma_i)$, which introduces a bias in the estimation.

Given our consistency result on $\hat{C}_{XX}^{(Q)}, \hat{C}_{XZ}^{(Q)}$, we can show that the estimated conditional operators are also consistent.

Theorem 3.6. *Suppose Assumptions 3.1 and 3.3. Given data $\{x_i, z_i\} \sim P$ and estimated covariance operators such that $\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \leq O_p(n^{-\alpha_2})$ and $\|\hat{C}_{XZ}^{(Q)} - C_{XZ}^{(Q)}\| \leq O_p(n^{-\alpha_2})$, by setting $\lambda = O(n^{-\frac{\alpha_2}{\beta_2+1}})$ we have*

$$\|\hat{E}_{Q,\lambda} - E_Q\| \leq O_p\left(n^{-\frac{\alpha_2\beta_2}{\beta_2+1}}\right).$$

The proof is in Appendix B.1. This rate is faster than the original KBR estimator (Fukumizu et al., 2013), which satisfies

$$\|\hat{E}_{Q,\lambda}^{(\text{org})} - E_Q\| \leq O_p\left(n^{-\frac{2\alpha_2\beta_2}{2\beta_2+5}}\right).$$

This is the benefit of avoiding the regularization of the form (2) and using instead (4). Given Corollary 3.5 and Theorem 3.6, we can thus show that

$$\|\hat{E}_{Q,\lambda} - E_Q\| \leq O_p\left(n^{-\frac{\alpha_1\beta_1\beta_2}{(\beta_1+1)(\beta_2+1)}}\right),$$

while Fukumizu et al. (2013) obtained

$$\|\hat{E}_{Q,\lambda}^{(\text{org})} - E_Q\| \leq O_p\left(n^{-\frac{\alpha_1(2\beta_1+1)2\beta_2}{(2\beta_1+2)(2\beta_2+5)}}\right).$$

The approach that yields the better overall rate depends on the smoothness parameters (β_1, β_2) . Our approach converges faster than the original KBR when the density ratio r_0 is smooth (i.e. $\beta_1 \simeq 1/2$) and the conditional operator

less smooth (i.e. $\beta_2 \simeq 0$). Note further that Sugiyama et al. (2008) show, even when $r_0 \neq \mathcal{H}_{\mathcal{Z}}$, that the KuLSIF estimator \hat{r} converges to the element in RKHS $\mathcal{H}_{\mathcal{Z}}$ with the least ℓ_2 error. We conjecture that our method might thus be robust to misspecification of the density estimator, although this remains a topic for future research (in particular, our proof requires consistency of the form in Theorem 3.4).

3.3. Learning Adaptive Features in KBR

Although kernel methods benefit from strong theoretical guarantees, a restriction to RKHS features limits our scope and flexibility, since this requires pre-specified feature maps. Empirically, poor performance can result in cases where the observable variables are high-dimensional (e.g. images), or have highly nonlinear relationships with the latents. Learned, adaptive neural network features have previously been used to substitute for kernel features when performing inference on mean embeddings in causal modeling (Xu et al., 2021a;b). Inspired by this work, we propose to employ adaptive NN observation features in KBR.

Recall that we learn the conditional operator $\hat{E}_{Q,\lambda}$ by minimizing the loss $\hat{\ell}^{(Q,\lambda)}$ defined in (3). We propose to jointly learn the feature map ψ_θ with the conditional operator $\hat{E}_{Q,\lambda}$ by minimizing the same loss,

$$\hat{\ell}^{(Q,\lambda)}(E, \theta) = \frac{1}{n} \sum_{i=1}^n \hat{r}_i \|\phi(z_i) - E^* \psi_\theta(x_i)\|^2 + \lambda \|E\|^2,$$

where $\psi_\theta : \mathcal{X} \rightarrow \mathbb{R}^d$ is the d -dimensional adaptive feature represented by a neural network parameterized by θ . As in the kernel feature case, the optimal operator can be obtained from (4) for a given value of θ . From this, we can write the loss for θ as

$$\begin{aligned} \ell(\theta) &= \arg \min_E \hat{\ell}^{(Q,\lambda)}(E, \theta) \\ &= \text{tr}\left(G_Z(D - D\Psi_\theta^\top(\Psi_\theta D\Psi_\theta^\top + \lambda I)^{-1}\Psi_\theta D)\right), \end{aligned}$$

where G_Z is the Gram matrix $(G_Z)_{ij} = k_{\mathcal{Z}}(z_i, z_j)$, D is the diagonal matrix with estimated importance weights $D = \text{diag}(\hat{r}_1, \dots, \hat{r}_n)$ and $\Psi_\theta \in \mathbb{R}^{n \times d}$ is the feature matrix defined as $\Psi_\theta = [\psi_\theta(x_1), \dots, \psi_\theta(x_n)]$. See Appendix B for the derivation.

The loss $\ell(\theta)$ can be minimized using gradient based optimization methods. Given the learned parameter $\hat{\theta} = \arg \min \ell(\theta)$ and the conditioning point \tilde{x} , we can estimate the posterior embedding $\hat{m}_{Q(Z|X)}(x) = \sum_{i=1}^n w_i \phi(z_i)$ with weights $\mathbf{w} = (w_1, \dots, w_n)$ given by

$$\mathbf{w} = D\Psi_{\hat{\theta}}^\top \left(\Psi_{\hat{\theta}} D\Psi_{\hat{\theta}}^\top + \lambda I \right)^{-1} \psi_{\hat{\theta}}(\tilde{x}).$$

Note that this corresponds to using a linear kernel on the learned features, $k_{\mathcal{X}}(x, x') = \psi_{\hat{\theta}}(x)^\top \psi_{\hat{\theta}}(x')$, in Algorithm 1. In experiments, we further employ a finite-dimensional random Fourier feature approximation of $\phi(z)$

to speed computation (Rahimi & Recht, 2008).

4. Kernel Bayes Filter

We next describe an important use-case for KBR, namely the *Kernel Bayes Filter (KBF)* (Fukumizu et al., 2013). Consider the following time invariant state-space model with observable variables X_t and hidden state variables Z_t ,

$$p(x_{1:T}, z_{1:T}) = p(x_1|z_1)p(z_1) \prod_{t=2}^T p(x_t|z_t)p(z_t|z_{t-1}).$$

Here, $X_{s:t}$ denotes $\{X_s, \dots, X_t\}$. Given this state-space model, the filtering problem aims to infer sequentially the distributions $p(z_t|x_1, \dots, x_t)$. Classically, filtering is solved by the Kalman filter, one of its nonlinear extensions (the extended Kalman filter (EKF) or unscented Kalman filter (UKF)), or a particle filter (Särkkä, 2013). These methods require knowledge of $p(x_t|z_t)$ and $p(z_t|z_{t-1})$, however. KBF does not require knowing these distributions and learns them from samples $(X_{1:T}, Z_{1:T})$ of both observable and hidden variables in a training phase.

Given a test sequence $\{\tilde{x}_1, \dots, \tilde{x}_T\}$, KBF sequentially applies KBR to obtain kernel embedding $m_{Z_t|\tilde{x}_{1:t}}$, where $m_{Z_t|\tilde{x}_{1:t}}$ denotes the embedding of the posterior distribution $P(z_t|X_{1:s} = \tilde{x}_{1:s})$. This can be obtained by iterating the following two steps. Assume that we have the embedding $m_{Z_t|\tilde{x}_{1:t}}$. Then, we can compute the embedding of forward prediction $m_{Z_{t+1}|\tilde{x}_{1:t}}$ by

$$m_{Z_{t+1}|\tilde{x}_{1:t}} = (E_{Z_{t+1}|Z_t})^* m_{Z_t|\tilde{x}_{1:t}},$$

where $E_{Z_{t+1}|Z_t}$ is the conditional operator for $P(z_{t+1}|z_t)$. Empirically, this is estimated from data $\{Z_t, Z_{t+1}\}$,

$$\hat{E}_{Z_{t+1}|Z_t} = (\hat{C}_{\text{prev,prev}} + \lambda'I)^{-1} \hat{C}_{\text{prev,post}},$$

where λ' is a regularizing coefficient and

$$\hat{C}_{\text{prev,prev}} = \frac{1}{T-1} \sum_{t=1}^{T-1} \phi(z_t) \otimes \phi(z_t),$$

$$\hat{C}_{\text{prev,post}} = \frac{1}{T-1} \sum_{t=1}^{T-1} \phi(z_t) \otimes \phi(z_{t+1}).$$

Given the probability of forward prediction $P(z_{t+1}|x_{1:t})$, we can obtain $P(z_{t+1}|x_{1:t+1})$ by applying Bayes' rule,

$$p(z_{t+1}|x_{1:t+1}) \propto p(x_{t+1}|z_{t+1})p(z_{t+1}|x_{1:t}).$$

Hence, we can obtain the filtering embedding $m_{Z_{t+1}|\tilde{x}_{1:t+1}}$ at the next timestep by applying KBR with prior embedding $m_{\Pi} = m_{Z_{t+1}|\tilde{x}_{1:t}}$ and samples $(X_{1:T}, Z_{1:T})$ at the conditioning point $\tilde{x} = \tilde{x}_{t+1}$. By repeating this process, we can conduct filtering for the entire sequence $\tilde{x}_{1:t}$.

Empirically, the estimated embedding $\hat{m}_{Z_t|\tilde{x}_{1:s}}$ is represented by a linear combination of features $\phi(z_t)$ as

Algorithm 2 Kernel Bayes Filter

Input: Training Sequence $\{(x_t, z_t)\}_{t=1}^T$, regularizing parameter $(\eta, \lambda, \lambda')$, Test sequence $\{\tilde{x}_t\}_{t=1}^{\tilde{T}}$.

- 1: Initialize $\mathbf{w}^{(0,1)}$
 - 2: **for** $t = 1$ **to** \tilde{T} **do**
 - 3: $\mathbf{w}^{(t,t)} \leftarrow$ Algorithm 1 for data $\{(x_t, z_t)\}_{t=1}^T$, prior embedding $\hat{m}_{\Pi} = \sum_{i=1} w_i^{(t-1,t)} \phi(z_i)$, conditioning point $\tilde{x} = \tilde{x}_t$.
 - 4: Compute $\mathbf{w}^{(t,t+1)}$ with $w_1^{(t,t+1)} = 0$ and $\mathbf{w}_{2:T}^{(t,t+1)} = (G_{Z_{-1}} + (T-1)\lambda'I)^{-1} \tilde{G}_{Z_{-1}} \mathbf{w}^{(t,t)}$ where $(G_{Z_{-1}})_{ij} = k_{\mathcal{Z}}(z_i, z_j) \in \mathbb{R}^{T-1 \times T-1}$ and $(\tilde{G}_{Z_{-1}})_{ij} = k_{\mathcal{Z}}(z_i, z_j) \in \mathbb{R}^{T-1 \times T-1}$.
 - 5: **end for**
-

$\hat{m}_{Z_t|\tilde{x}_{1:s}} = \sum_{i=1} w_i^{(s,t)} \phi(z_i)$. The update equations for weights $w_i^{(s,t)}$ are summarized in Algorithm 2. If we have any prior knowledge of $P(z_1)$, we can initialize $\mathbf{w}^{(0,1)}$ accordingly. Otherwise, we can set $\mathbf{w}^{(0,1)} = (1/T, \dots, 1/T)^\top$, which initialization is used in the experiments.

5. Related Work on Neural Filtering

Several recent methods have been proposed combining state-space models with neural networks (Krishnan et al., 2015; Klushyn et al., 2021; Rangapuram et al., 2018), aiming to learn the latent dynamic and the observation models from observed sequences $X_{1:T}$ alone. These approaches assume a *parametric* form of the latent dynamics: for example, the Deep Kalman filter (Krishnan et al., 2015) assumes that the distribution of the latents is Gaussian, with mean and covariance which are nonlinear functions of the previous latent state. DeepSSM (Rangapuram et al., 2018) assumes linear latent dynamics with Gaussian noise, and EKVAE (Klushyn et al., 2021) uses a locally linear Gaussian transition model. These models use the variational inference techniques to learn the parameters, which makes it challenging to prove the convergence to the true models. In contrast, KBF learns latent dynamics and observation model nonparametrically from samples, and the accuracy of the filtering is guaranteed from the convergence of KBR.

6. Experiments

In this section, we empirically investigate the performance of our KBR estimator in a variety of settings, including the problem of learning posterior mean proposed in Fukumizu et al. (2013), as well as challenging filtering problems where the observations are high-dimensional images.

6.1. Learning Posterior Mean From Samples

We revisit here the problem introduced by Fukumizu et al. (2013), which learns the posterior mean from samples. Let $X \in \mathbb{R}^d$ and $Z \in \mathbb{R}^d$ be generated from $P(X, Z) = \mathcal{N}((\mathbf{1}_d^\top, \mathbf{0}_d^\top)^\top, V)$, where $\mathbf{1}_d = (1, \dots, 1)^\top \in \mathbb{R}^d$, $\mathbf{0}_d = (0, \dots, 0)^\top \in \mathbb{R}^d$ and $V = \frac{1}{2d}A^\top A + 2I$, with each component of $A \in \mathbb{R}^{2d \times 2d}$ sampled from the standard Gaussian distribution on each run. The prior is set to be $\Pi(Z) = \mathcal{N}(\mathbf{0}_d, V_{ZZ}/2)$, where V_{ZZ} is the Z -component of V . We construct the prior embedding using 200 samples from $\Pi(Z)$, and we aim to predict the mean of posterior $\mathbb{E}_Q[Z|X]$ using $n = 200$ data points from $P(X, Z)$, as the dimension d varies.

Figure 1 summarizes the MSEs over 30 runs when the conditioning points \tilde{x} are sampled from $\mathcal{N}(\mathbf{0}_d, V_{XX})$. Here, Original denotes the performance of the original KBR estimator using operator $\hat{E}_{Q,\lambda}^{(\text{org})}$ in (2), and IW is computed from IW-KBR approach, which uses operator $\hat{E}_{Q,\lambda}$ in (4). In this experiment, we set $\eta = \lambda = 0.2$ and used Gaussian kernels for both KBR methods, where the bandwidth is given by the median trick. Unsurprisingly, the error increases as dimension increases. However, the IW-KBR estimator performs significantly better than the original KBR estimator. This illustrates the robustness of the IW-KBR approach even when the model is misspecified, since the correct $\mathbb{E}_Q[Z|X]$, which is a linear function of X , does not belong to the Gaussian RKHS.

To show how the quality of the density ratio estimate relates to the overall performance, we include the result of IWKBR with the true density ratio r_0 , denoted as “IW-TrueRatio” in Figure 1. Surprisingly, IW-TrueRatio performs worse than the original IWKBR, which estimates density ratio using KuLSIF. This suggests that true density ratio yields a suboptimal bias/variance tradeoff, and greater variance for the *finite sample estimate of the KBR posterior*, compared with the smoothed estimate obtained from KuLSIF.

6.2. Low-Dimensional KBF

We consider a filtering problem introduced by Fukumizu et al. (2013), with latent $Z_t \in \mathbb{R}^2$ and observation $X_t \in \mathbb{R}^2$. The dynamics of the latent $Z_t = (u_t, v_t)^\top$ are given as follows. Let $\theta_t \in [0, 2\pi]$ be

$$\cos \theta_t = \frac{u_t}{\sqrt{u_t^2 + v_t^2}}, \quad \sin \theta_t = \frac{v_t}{\sqrt{u_t^2 + v_t^2}}.$$

The latent $Z_{t+1} = (u_{t+1}, v_{t+1})^\top$ is then

$$\begin{pmatrix} u_{t+1} \\ v_{t+1} \end{pmatrix} = (1 + \beta \sin(M\theta_t)) \begin{pmatrix} \cos(\theta_t + \omega) \\ \sin(\theta_t + \omega) \end{pmatrix} + \epsilon_Z, \quad (5)$$

for given parameters (β, M, ω) . The observation X_t is given by $X_t = Z_t + \epsilon_X$. Here, ϵ_X and ϵ_Z are noise variables sampled from $\epsilon_X \sim \mathcal{N}(0, \sigma_X^2 I)$ and $\epsilon_Z \sim \mathcal{N}(0, \sigma_Z^2 I)$.

Fukumizu et al. test performance using the “Rotation” dynamics $\omega = 0.3, b = 0$ and “Oscillatory” dynamics $\omega = 0.4, b = 0.4, M = 8$, where the noise level is set to $\sigma_X = \sigma_Z = 0.2$ in both scenarios. Using the same dynamics, we evaluate the performance of our proposed estimator by the MSE in predicting Z_t from $X_{1:t}$, where the length of the test sequence is set to 200. We repeated the experiments 30 times for each setting.

Results are summarized in Figure 2: Original denotes the results for using original KBR approach in Algorithm 2, while IW used our IW-KBR approach. For both approaches, we used Gaussian kernels k_X, k_Z whose bandwidths are set to βD_X and βD_Z , respectively. Here, D_X and D_Z are the medians of pairwise distances among the training samples. We used the KuLSIF leave-one-out cross-validation procedure (Kanamori et al., 2012) to tune the regularization parameter η , and set $\lambda' = \eta$. This leaves two parameters to be tuned: the scaling parameter β and the regularization parameter λ . These are selected using the last 200 steps of the training sequence as a validation set. We also include the results for the extended Kalman filter (EKF) and the particle filter (Särkkä, 2013).

Figure 2 shows that the EKF and the particle filter perform slightly better than KBF methods in “Rotation” dynamics, which replicates the results in Fukumizu et al.. This is not surprising, since these methods have access to the true dynamics, which makes the tracking easier. In “Oscillatory” dynamics, however, which has a stronger nonlinearity, KBF displays comparable or better performance than the EKF, which suffers from a large error caused by the linear approximation. In both scenarios, IW-KBR slightly outperforms the original KBR.

6.3. High-Dimensional KBF

Finally, we apply KBF to scenarios where observations are given by images. We set up two experiments: one uses high dimensional complex images while the latent follows simple dynamics, while the other considers complex dynamics with observations given by relatively simple images. In both cases, the particle filter performs significantly worse than the methods based on neural networks, since likelihood evaluation is unstable in the high dimensional observation space, despite the true dynamics being available.

Deepmind Lab Video: The first high-dimensional KBF experiment uses DeepMind Lab (Beattie et al., 2016), which is a 3D game-like platform for reinforcement learning. This platform provides a first-person viewpoint through the eyes of the agent. An example image can be found in Figure 3. Based on this platform, we introduce a filtering problem to estimate the agent’s orientation at a specific point in the maze.




Figure 1. MSE as a function of dimension for posterior mean prediction




Figure 2. Results for low-dimensional KBF experiments. The asterisks indicates the level of statistical significance (*: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$, ****: $p < 0.0001$), which is tested by the Wilcoxon signed-rank test.




Figure 3. DeepMind Lab




Figure 4. dSprite

The dynamics are as follows. Let $\theta_t \in [0, 2\pi)$ be the direction that an agent is facing at time t . The next direction is

$$\theta_{t+1} = \theta_t + \omega + \epsilon_\theta \pmod{2\pi},$$

where noise $\epsilon_\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. Let $Z_t = (\cos \theta_t, \sin \theta_t)$ and X_t be the image that corresponds to the direction $\theta_t + \epsilon_X$ where $\epsilon_X \sim \mathcal{N}(0, \sigma_X^2)$. We used the RGB images down-scaled to $(36, 48)$, which makes the observation $36 \times 48 \times 3 = 5184$ dimensions. We added Gaussian noise $\mathcal{N}(0, 0.01)$ to the each dimension of the observations Z_t . We ran the experiments with $\omega = 0.4, \sigma_\theta = 0.2, \sigma_X = 0.2$, MSEs in 30 runs are summarized in Figure 5.




Figure 5. Result for DeepMind Lab setting

In addition to `Original` and `IW`, whose hyper-parameters are selected by the same procedure as in the low-dimensional experiments, we introduce two neural network based methods: `IW(NN)` and `LSTM`. `IW(NN)` uses adaptive feature discussed in Section 3.3. Here, instead of learning different features for each timestep, which would be time consuming and redundant, we learn adaptive feature ψ_θ by minimizing

$$\hat{\ell}^{(P, \lambda)}(E, \theta) = \frac{1}{n} \sum_{i=1}^n \|\tilde{\phi}(z_i) - E^* \psi_\theta(x_i)\|^2 + \lambda \|E\|^2$$

with the learned parameter $\hat{\theta}$ being used for all timesteps. `LSTM` denotes an LSTM baseline (Hochreiter & Schmidhuber, 1997), which predicts Z_t from features extracted from $X_{t-10:t}$. To make the comparison fair, `LSTM` used the same network architecture in the feature extractor as `IW(NN)`.

As in low-dimensional cases, `IW` consistently performs better for all sequence lengths (Figure 5). The baseline `LSTM` performs similarly to `IW`, however, even though it does not explicitly model the dynamics. This is because functions in the RKHS are not expressive enough to model the relationship between the direction and the images. This is mitigated by using adaptive features in `IW(NN)`, which outperforms other methods by taking the advantage of the strong expressive power of neural networks.

dSprite Setting: The second high-dimensional KBF experiment uses `dSprite` (Matthey et al., 2017), which is a dataset of 2D shape images as shown in Figure 4. Here, we consider the latent Z_t following the same dynamics as (5), and the model observes the image X_t where the position of the shape corresponds to the noisy observation of the latent $Z_t + \epsilon_X, \epsilon_X \sim \mathcal{N}(0, \sigma_X^2 I)$. We used the ‘‘Oscillatory’’ dynamics (i.e. $\omega = 0.4, b = 0.4, M = 8$) and set noise levels to $\sigma_X = 0.05, \sigma_Z = 0.2$. Again, we added Gaussian noise $\mathcal{N}(0, 0.01)$ to the images.

MSEs across 30 runs are summarized in Figure 6. Unlike in DeepMind Lab setting, `LSTM` performs the worst in this setting. This suggests the advantage of KBF methods, which explicitly model the dynamics and exploit them in the filtering. Among KBF methods, `IW` and `IW(NN)` perform significantly better than `Original`, demonstrating the superiority of the `IW` approach.

7. Conclusion

In this paper, we proposed a novel approach to kernel Bayes’ rule, `IW-KBR`, which minimizes a loss estimated by impor-




Figure 6. Result for dSprite setting

tance weighting. We established consistency of IW-KBR based on a novel analysis of an RKHS density ratio estimate, which is of independent interest. Our empirical evaluation shows that the IW-KBR significantly outperforms the existing estimator in both Bayesian inference and filtering for state-space models. Furthermore, by learning adaptive neural net features, IW-KBR outperforms a neural sequence model in filtering problems with high-dimensional image observations.

In future work, we suggest exploring different density ratio estimation techniques for our setting. It is well-known in the density ratio estimation context that KuLSIF estimator may suffer from high variance. To mitigate this, Yamada et al. (2011) proposed to use relative density-ratio estimation. Deriving a consistency result for such an estimator and applying it in kernel Bayes' rule would be a promising approach. It would further be of interest to apply IW-KBR in additional settings, such as approximate Bayesian computation (Tavaré et al., 1997; Marjoram et al., 2003), as also discussed by Fukumizu et al. (2013).

Acknowledgements

This work was supported by the Gatsby Charitable Foundation.

References

Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., Schrittwieser, J., Anderson, K., York, S., Cant, M., Cain, A., Bolton, A., Gaffney, S., King, H., Hassabis, D., Legg, S., and Petersen, S. Deepmind lab, 2016.

Boots, B., Gordon, G., and Gretton, A. Hilbert space embeddings of predictive state representations. In *Uncertainty in Artificial Intelligence*, 2013.

Caponnetto, A. and Vito, E. D. Optimal rates for the regu-

larized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.

Chwialkowski, K., Strathmann, H., and Gretton, A. A kernel test of goodness of fit. In *International Conference on Machine Learning*, 2016.

Cucker, F. and Smale, S. On the mathematical foundations of learning. *Bulletin of the American Mathematical Society*, 39:1–49, 2001.

Fukumizu, K., Bach, F. R., and Jordan, M. I. Kernel dimension reduction in regression. *The Annals of Statistics*, 37(4):1871–1905, 2009.

Fukumizu, K., Song, L., and Gretton, A. Kernel Bayes' rule: Bayesian inference with positive definite kernels. *Journal of Machine Learning Research*, 14(82):3753–3783, 2013.

Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.

Grunewalder, S., Lever, G., Baldassarre, L., Patterson, S., Gretton, A., and Pontil, M. Conditional mean embeddings as regressors. In *International Conference on Machine Learning*, 2012.

Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

Jasra, A., Singh, S. S., Martin, J. S., and McCoy, E. Filtering via approximate Bayesian computation. *Statistics and Computing*, 22(6):1223–1237, 2012.

Jitkrittum, W., Gretton, A., Heess, N., Eslami, S., Lakshminarayanan, B., Sejdinovic, D., and Szabo, Z. Kernel-based just-in-time learning for passing expectation propagation messages. In *Uncertainty in Artificial Intelligence*, 2015.

Kanagawa, M., Nishiyama, Y., Gretton, A., and Fukumizu, K. Filtering with state-observation examples via kernel Monte Carlo filter. *Neural Computation*, 28(2):382–444, 2016.

Kanamori, T., Hido, S., and Sugiyama, M. A least-squares approach to direct importance estimation. *Journal of Machine Learning Research*, 10(48):1391–1445, 2009.

Kanamori, T., Suzuki, T., and Sugiyama, M. Statistical analysis of kernel-based least-squares density-ratio estimation. *Machine Learning*, 86(3):335–367, 2012.

Klushyn, A., Kurle, R., Soelch, M., Cseke, B., and van der Smagt, P. Latent matters: Learning deep state-space models. In *Advances in Neural Information Processing Systems*, volume 34, pp. 10234–10245, 2021.

-
- Krishnan, R. G., Shalit, U., and Sontag, D. Deep kalman filters, 2015.
- Marjoram, P., Molitor, J., Plagnol, V., and Tavaré, S. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 100(26):15324–15328, 2003.
- Matthey, L., Higgins, I., Hassabis, D., and Lerchner, A. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- Muandet, K., Fukumizu, K., Dinuzzo, F., and Schölkopf, B. Learning from distributions via support measure machines. In *Advances in Neural Information Processing Systems*, 2012.
- Muandet, K., Kanagawa, M., Saengkyongam, S., and Marukatat, S. Counterfactual mean embeddings. *Journal of Machine Learning Research*, 22(162):1–71, 2021.
- Nishiyama, Y., Kanagawa, M., Gretton, A., and Fukumizu, K. Model-based kernel sum rule. *Machine Learning*, 109: 939–972, 2020.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, volume 20, 2008.
- Rangapuram, S. S., Seeger, M. W., Gasthaus, J., Stella, L., Wang, Y., and Januschowski, T. Deep state space models for time series forecasting. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- Särkkä, S. *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013.
- Singh, R., Sahani, M., and Gretton, A. Kernel instrumental variable regression. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Smale, S. and Zhou, D.-X. Learning theory estimates via integral operators and their approximations. *Constructive Approximation*, 26(2):153–172, 2007.
- Smola, A. J., Gretton, A., Song, L., and Schölkopf, B. A Hilbert space embedding for distributions. In *Algorithmic Learning Theory*, volume 4754, pp. 13–31. Springer, 2007.
- Song, L., Huang, J., Smola, A., and Fukumizu, K. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *International Conference on Machine Learning*, 2009.
- Sriperumbudur, B. K., Gretton, A., Fukumizu, K., Schölkopf, B., and Lanckriet, G. R. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.
- Sugiyama, M., Suzuki, T., Nakajima, S., Kashima, H., von Bünau, P., and Kawanabe, M. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60:699–746, 2008.
- Szabo, Z., Sriperumbudur, B., Poczos, B., and Gretton, A. Learning theory for distribution regression. *Journal of Machine Learning Research*, 17(152):1–40, 2016.
- Tavaré, S., Balding, D. J., Griffiths, R. C., and Donnelly, P. Inferring coalescence times from dna sequence data. *Genetics*, 145(2):505–518, 1997.
- Williams, C. and Seeger, M. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, volume 13, 2001.
- Xu, L., Chen, Y., Srinivasan, S., de Freitas, N., Doucet, A., and Gretton, A. Learning deep features in instrumental variable regression. In *International Conference on Learning Representations*, 2021a.
- Xu, L., Kanagawa, H., and Gretton, A. Deep proxy causal learning and its application to confounded bandit policy evaluation. In *Advances in Neural Information Processing Systems*, 2021b.
- Yamada, M., Suzuki, T., Kanamori, T., Hachiya, H., and Sugiyama, M. Relative density-ratio estimation for robust distribution comparison. In *Advances in Neural Information Processing Systems*, 2011.

A. Implication of Assumption 3.2

In this appendix, we will discuss the implication of Assumption 3.2 when we consider the ratio of two Gaussian distributions. Let $\pi \sim \mathcal{N}(0, 1)$ and $p \sim \mathcal{N}(0, 2)$. Then, the density ratio is

$$r_0(z) = \frac{\pi(z)}{p(z)} = \sqrt{2} \exp\left(-\frac{z^2}{4}\right),$$

which is in the RKHS $\mathcal{H}_{\mathcal{Z}}$ induced by Gaussian kernel $k(z, z') = \exp(-(z - z')^2/4)$. Indeed, we can see that

$$r_0(z) = \sqrt{2}k(z, 0).$$

Note that, from reproducing property, we have $\langle r_0, f \rangle = \sqrt{2}f(0)$ for all $f \in \mathcal{H}_{\mathcal{Z}}$.

Given this, we can analytically compute the eigendecomposition of the covariance operator as

$$C_{ZZ}^{(P)} = \sum_i \lambda_i (\sqrt{\lambda_i} e_i(z)) \otimes (\sqrt{\lambda_i} e_i(z)),$$

where $\lambda_i = O(B^k)$ with constant $B < 1$ and e_i is the i -th order Hermite polynomial (Rasmussen & Williams, 2006, Section 4.3). Assumption 3.2 requires $\|(C_{ZZ}^{(P)})^{-\beta_1} r_0\|^2$ finite, meaning β_1 is the maximum value that satisfies

$$\|(C_{ZZ}^{(P)})^{-\beta_1} r_0\|^2 = \sum_i \lambda_i^{-2\beta_1} \langle \sqrt{\lambda_i} e_i, r_0 \rangle^2 = 2 \sum_i \lambda_i^{1-2\beta_1} e_i^2(0) \leq \infty.$$

B. Proof of Theorems

In this appendix, we provide the proof of our theorems.

B.1. Proof for Convergence Analysis

We will rely on the following concentration inequality.

Proposition B.1 (Lemma 2 of (Smale & Zhou, 2007)). *Let ξ be a random variable taking values in a real separable RKHS \mathcal{H} with $\|\xi\| \leq M$ almost surely, and let ξ_1, \dots, ξ_n be i.i.d. random variables distributed as ξ . Then, for all $n \in \mathbb{N}$ and $\delta \in (0, 1)$,*

$$\mathbb{P} \left[\left\| \frac{1}{n} \sum_{i=1}^n \xi_i - \mathbb{E}[\xi] \right\| \leq \frac{2M \log 2/\delta}{n} + \sqrt{\frac{2\mathbb{E}[\|\xi\|^2] \log 2/\delta}{n}} \right] \geq 1 - \delta.$$

B.1.1. PROOF OF THEOREM 3.4 AND COROLLARY 3.5

We review some properties of the density ratio r_0 used in KuLSIF.

Lemma B.2 ((Kanamori et al., 2009)). *If the density ratio $r_0 \in \mathcal{H}_{\mathcal{Z}}$, then we have*

$$C_{ZZ}^{(P)} r_0 = m_{\Pi},$$

where $m_{\Pi} = \mathbb{E}_{\Pi}[\phi(Z)]$.

Proof. From reproducing characteristics, we have

$$C_{ZZ}^{(P)} r_0 = \mathbb{E}_P[r_0(Z)\phi(Z)] = \mathbb{E}_{\Pi}[\phi(Z)].$$

The last equality holds from the definition of the density ratio $r_0(z) = \pi(z)/p(z)$. \square

Given Lemma B.2, we can bound the error of untruncated KuLISF estimator in RKHS norm. Let \tilde{r}_{η} be

$$\tilde{r}_{\eta} = \left(\hat{C}_{ZZ}^{(P)} + \eta I \right)^{-1} \hat{m}_{\Pi}.$$

Note that the weight γ_i appearing in the original KBR (1) can be understood as the value of \tilde{r}_{η} at specific data points: $\gamma_i = \tilde{r}_{\eta}(z_i)$. Hence, the weight we use can be written as

$$\hat{r}_i = \max(0, \tilde{r}_{\eta}(z_i)).$$

Although we use truncated weights as above in KBR, we first derive an upper-bound on the error $\|r_0 - \tilde{r}_\eta\|$. Furthermore, we define the function r_η which is a popular version of \tilde{r}_η .

$$r_\eta = \left(C_{ZZ}^{(P)} + \eta I \right)^{-1} m_\Pi.$$

Using r_η , we can decompose the error as

$$\|\tilde{r}_\eta - r_0\| \leq \|\tilde{r}_\eta - r_\eta\| + \|r_\eta - r_0\|.$$

The second term can be bounded using an approach similar to Theorem 6 in Singh et al. (2019).

Lemma B.3. *Suppose Assumption 3.2, then we have*

$$\|r_\eta - r_0\| \leq \eta^{\beta_1} \zeta_1.$$

Proof. Let (ν_k, e_k) be the eigenvalues and eigenfunctions of $C_{ZZ}^{(P)}$, then

$$\begin{aligned} \|r_\eta - r_0\|^2 &= \|((C_{ZZ}^{(P)} + \eta I)^{-1} C_{ZZ}^{(P)} - I)r_0\|^2 \\ &= \sum_k \left(\frac{\nu_k}{\nu_k + \eta} - 1 \right)^2 \langle r_0, e_k \rangle^2 \\ &= \sum_k \left(\frac{\eta}{\nu_k + \eta} \right)^2 \langle r_0, e_k \rangle^2 \left(\frac{\eta}{\eta} \cdot \frac{\nu_k}{\nu_k} \cdot \frac{\eta + \nu_k}{\eta + \nu_k} \right)^{2\beta_1} \\ &= \eta^{2\beta_1} \sum_k \nu_k^{-2\beta_1} \langle r_0, e_k \rangle^2 \left(\frac{\eta}{\nu_k + \eta} \right)^{2-2\beta_1} \left(\frac{\nu_k}{\nu_k + \eta} \right)^{2\beta_1} \\ &\leq \eta^{2\beta_1} \sum_k \nu_k^{-2\beta_1} \langle r_0, e_k \rangle^2 = \eta^{2\beta_1} \zeta_1^2, \end{aligned}$$

where the first equality uses Lemma B.2, and the last equality holds from Assumption 3.2 and the fact that

$$\|g_1\|^2 = \|(C_{ZZ}^{(P)})^{-\beta_1} r_0\|^2 = \sum_k \nu_k^{-2\beta_1} \langle r_0, e_k \rangle^2.$$

□

We now establish a useful lemma on the norm of functions.

Lemma B.4. $\|r_\eta\| \leq \|r_0\|$.

Proof. Let (ν_k, e_k) be the eigenvalues and eigenfunctions of $C_{ZZ}^{(P)}$. From Lemma B.2, we have

$$\begin{aligned} \|r_\eta\|^2 &= \|((C_{ZZ}^{(P)} + \eta I)^{-1} C_{ZZ}^{(P)})r_0\|^2 \\ &= \sum_k \left(\frac{\nu_k}{\nu_k + \eta} \right)^2 \langle r_0, e_k \rangle^2 \\ &\leq \sum_k \langle r_0, e_k \rangle^2 = \|r_0\|^2. \end{aligned}$$

□

Now, we are ready to prove our convergence result on \tilde{r}_η .

Theorem B.5. *Suppose Assumptions 3.1 and 3.2. Given data $\{x_i, z_i\} \sim P$ and the estimated prior embedding \hat{m}_Π , we have*

$$\|\tilde{r}_\eta - r_0\| \leq \frac{1}{\eta} \left(\frac{4\kappa_1^2 \|r_0\| \log 2/\delta}{\sqrt{n}} + \|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \right) + \eta^{\beta_1} \zeta_1$$

with probability at least $1 - \delta$ for $\delta \in (0, 2/e)$.

Proof. We decompose the error as

$$\|\tilde{r}_\eta - r_0\| \leq \|\tilde{r}_\eta - r_\eta\| + \|r_\eta - r_0\|.$$

From Lemma B.3, we have $\|r_\eta - r_0\| \leq \eta^{\beta_1} \zeta_1$. For the first term, we have

$$\begin{aligned} \|r_\eta - \tilde{r}_\eta\| &= \|(\hat{C}_{ZZ}^{(P)} + \eta I)^{-1} (\hat{C}_{ZZ}^{(P)} r_\eta + \eta r_\eta - \hat{m}_\Pi)\| \\ &\leq \|(\hat{C}_{ZZ}^{(P)} + \eta I)^{-1}\| \| \hat{C}_{ZZ}^{(P)} r_\eta + (m_\Pi - C_{ZZ}^{(P)} r_\eta) - \hat{m}_\Pi \| \\ &\leq \frac{1}{\eta} \left(\|(\hat{C}_{ZZ}^{(P)} - C_{ZZ}^{(P)}) r_\eta\| + \|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \right). \end{aligned}$$

By applying Proposition B.1 with $\xi = (\phi(Z) \otimes \phi(Z)) r_\eta$, we have

$$\|(\hat{C}_{ZZ}^{(P)} - C_{ZZ}^{(P)}) r_\eta\| \leq \frac{2\kappa_1^2 \|r_0\| \log(2/\delta)}{n} + \sqrt{\frac{2\kappa_1^4 \|r_0\|^2 \log(2/\delta)}{n}}$$

with probability $1 - \delta$ since

$$\begin{aligned} \|\xi\| &\leq \|\phi(Z) \otimes \phi(Z)\| \|r_\eta\| \leq \kappa_1^2 \|r_\eta\| \leq \kappa_1^2 \|r_0\|, \\ \mathbb{E}[\|\xi\|^2] &\leq \kappa_1^4 \|r_\eta\|^2 \leq \kappa_1^4 \|r_0\|^2 \end{aligned}$$

from Lemma B.4. Since $2\log(2/\delta) \geq \sqrt{2}\log(2/\delta)$ for $\delta \in (0, 2/e)$, we finally obtain

$$\|\tilde{r}_\eta - r_0\| \leq \frac{1}{\eta} \left(\frac{4\kappa_1^2 \|r_0\| \log 2/\delta}{\sqrt{n}} + \|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \right) + \eta^{\beta_1} \zeta_1.$$

□

Given Theorem B.5, Theorem 3.4 is now easy to prove.

Proof of Theorem 3.4. Note that we have

$$\begin{aligned} \|\hat{r} - r_0(z)\|_\infty &= \max_{z \in \mathcal{Z}} |\max(\tilde{r}_\eta(z), 0) - r_0(z)| \\ &\leq \max_{z \in \mathcal{Z}} |\tilde{r}_\eta(z) - r_0(z)| \\ &\leq \|\tilde{r}_\eta - r_0\| \max_{z \in \mathcal{Z}} \|\phi(z)\| \\ &\leq \kappa_1 \|\tilde{r}_\eta - r_0\|, \end{aligned}$$

where the first inequality uses the fact that density ratio r_0 is non-negative. From the upper-bound in Theorem B.5, we obtain

$$\begin{aligned} \|\hat{r} - r_0(z)\|_\infty &\leq \kappa_1 \|\tilde{r}_\eta - r_0\| \\ &\leq \frac{\kappa_1}{\eta} \left(\frac{4\kappa_1^2 \|r_0\| \log 2/\delta}{\sqrt{n}} + \|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \right) + \eta^{\beta_1} \kappa_1 \zeta_1. \end{aligned}$$

with probability $1 - \delta$ for $\delta \in (0, 2/e)$. Hence, if $\|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \leq O_p(n^{-\alpha_1})$, by setting $\eta = O(n^{-\frac{\alpha_1}{\beta_1+1}})$, we have

$$\|\hat{r} - r_0(z)\|_\infty \leq O_p(n^{-\frac{\alpha_1 \beta_1}{\beta_1+1}}).$$

□

Furthermore, we can show the following convergence result on the covariance operator as follows.

Proof of Corollary 3.5. From Theorem 3.4, with probability $1 - \delta$ for $\delta \in (0, 2/e)$, we have

$$\max_i |r_i - r_0(z_i)| \leq \|\hat{r} - r_0\|_\infty \leq \frac{\kappa_1}{\eta} \left(\frac{4\kappa_1^2 \|r_0\| \log 2/\delta}{\sqrt{n}} + \|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \right) + \eta^{\beta_1} \kappa_1 \zeta_1.$$

Let $\tilde{C}_{XX}^{(Q)}$ be the empirical covariance operator with true density ratio r_0 :

$$\tilde{C}_{XX}^{(Q)} = \sum_{i=1}^n r_0(z_i) \psi(x_i) \otimes \psi(x_i).$$

Then, we have

$$\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \leq \|\hat{C}_{XX}^{(Q)} - \tilde{C}_{XX}^{(Q)}\| + \|C_{XX}^{(Q)} - \tilde{C}_{XX}^{(Q)}\|.$$

For the first term, we have

$$\begin{aligned} \|\hat{C}_{XX}^{(Q)} - \tilde{C}_{XX}^{(Q)}\| &\leq \frac{1}{n} \sum_i |r_i - r_0(z_i)| \|\psi(x_i) \otimes \psi(x_i)\| \\ &\leq \frac{\kappa_1 \kappa_2^2}{\eta} \left(\frac{4\kappa_1^2 \|r_0\| \log 2/\delta}{\sqrt{n}} + \|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \right) + \eta^{\beta_1} \kappa_1 \kappa_2^2 \zeta_1. \end{aligned}$$

For the second term, by applying Proposition B.1 with $\xi = r_0(Z)(\psi(X) \otimes \psi(X))$, we have

$$\begin{aligned} \|C_{XX}^{(Q)} - \tilde{C}_{XX}^{(Q)}\| &\leq \frac{2\kappa_1 \kappa_2^2 \|r_0\| \log 2/\delta}{n} + \sqrt{\frac{2(\kappa_1 \kappa_2^2 \|r_0\|)^2 \log 2/\delta}{n}} \\ &\leq \frac{4\kappa_1 \kappa_2^2 \|r_0\| \log 2/\delta}{\sqrt{n}} \end{aligned}$$

with probability $1 - \delta$ for $\delta \in (0, 2/e)$ since

$$\begin{aligned} \|\xi\| &= |r_0| \|\psi(X) \otimes \psi(X)\| \\ &\leq \kappa_1 \kappa_2^2 \|r_0\|, \\ \mathbb{E}[\|\xi\|^2] &\leq (\kappa_1 \kappa_2^2 \|r_0\|)^2. \end{aligned}$$

Hence, we have

$$\begin{aligned} \|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| &\leq \|\hat{C}_{XX}^{(Q)} - \tilde{C}_{XX}^{(Q)}\| + \|C_{XX}^{(Q)} - \tilde{C}_{XX}^{(Q)}\| \\ &\leq \frac{\kappa_1 \kappa_2^2}{\eta} \left(\frac{4\kappa_1^2 \|r_0\| \log 2/\delta}{\sqrt{n}} + \|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \right) + \eta^{\beta_1} \kappa_1 \kappa_2^2 \zeta_1 + \frac{4\kappa_1 \kappa_2^2 \|r_0\| \log 2/\delta}{\sqrt{n}} \end{aligned}$$

with probability $1 - 2\delta$ for $\delta \in (0, 2/e)$. Thus, if $\|\hat{m}_\Pi - \mathbb{E}_\Pi[\phi(X)]\| \leq O_p(n^{-\alpha_1})$, by setting $\eta = O(n^{-\frac{\alpha_1}{\beta_1+1}})$, we have

$$\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \leq O_p(n^{-\frac{\alpha_1 \beta_1}{\beta_1+1}}).$$

□

B.1.2. PROOF OF THEOREM 3.6

Next, we derive Theorem 3.6. Again, we define the operator $E_{Q,\lambda}$ which replaces the empirical estimates in $\hat{E}_{Q,\lambda}$ by their population versions.

$$E_{Q,\lambda} = (C_{XX}^{(Q)} + \lambda I)^{-1} C_{XX}^{(Q)}$$

By following a similar approach as in Lemma B.3, we obtain the following result.

Lemma B.6 (Theorem 6 in (Singh et al., 2019)). *Suppose Assumption 3.3, then we have*

$$\|E_{Q,\lambda} - E_Q\| \leq \lambda^{\beta_2} \zeta_2.$$

Using a proof similar to the one of Lemma B.4, we also have

Lemma B.7. $\|E_{Q,\lambda}\| \leq \|E_Q\|.$

Given these lemmas, we can prove Theorem 3.6 as follows.

Proof of Theorem 3.6. We decompose the error as follows:

$$\|E_Q - \hat{E}_{Q,\lambda}\| \leq \|E_Q - E_{Q,\lambda}\| + \|E_{Q,\lambda} - \hat{E}_{Q,\lambda}\|.$$

Lemma B.6 bounds the first term as

$$\|E_{Q,\lambda} - E_Q\| \leq \lambda^{\beta_2} \zeta_2.$$

The second term can be bounded as follows

$$\begin{aligned} \|E_{Q,\lambda} - \hat{E}_{Q,\lambda}\| &= \|(\hat{C}_{XX}^{(Q)} + \lambda I)^{-1}(\hat{C}_{XX}^{(Q)} E_{Q,\lambda} + \lambda E_{Q,\lambda} - \hat{C}_{XZ}^{(Q)})\| \\ &\leq \|(\hat{C}_{XX}^{(Q)} + \lambda I)^{-1}\| \|(\hat{C}_{XX}^{(Q)} E_{Q,\lambda} + \lambda E_{Q,\lambda} - \hat{C}_{XZ}^{(Q)})\| \\ &\leq \frac{1}{\lambda} \|(\hat{C}_{XX}^{(Q)} E_{Q,\lambda} + \lambda E_{Q,\lambda} - \hat{C}_{XZ}^{(Q)})\| \\ &= \frac{1}{\lambda} \|(\hat{C}_{XX}^{(Q)} E_{Q,\lambda} + (C_{XZ}^{(Q)} - C_{XX}^{(Q)} E_{Q,\lambda}) - \hat{C}_{XZ}^{(Q)})\| \\ &= \frac{1}{\lambda} \|(\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}) E_{Q,\lambda} - (\hat{C}_{XZ}^{(Q)} - C_{XZ}^{(Q)})\| \\ &= \frac{1}{\lambda} (\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \|E_{Q,\lambda}\| + \|\hat{C}_{XZ}^{(Q)} - C_{XZ}^{(Q)}\|) \\ &= \frac{1}{\lambda} (\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \|E_Q\| + \|\hat{C}_{XZ}^{(Q)} - C_{XZ}^{(Q)}\|) \end{aligned}$$

where the last inequality uses Lemma B.7. Therefore, we have

$$\|E_Q - \hat{E}_{Q,\lambda}\| \leq \frac{1}{\lambda} (\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \|E_Q\| + \|\hat{C}_{XZ}^{(Q)} - C_{XZ}^{(Q)}\|) + \lambda^{\beta_2} \zeta_2.$$

Hence, if $\|\hat{C}_{XX}^{(Q)} - C_{XX}^{(Q)}\| \leq O_p(n^{-\alpha_2})$ and $\|\hat{C}_{XZ}^{(Q)} - C_{XZ}^{(Q)}\| \leq O_p(n^{-\alpha_2})$, we can see

$$\|E_Q - \hat{E}_{Q,\lambda}\| \leq O_p(n^{-\frac{\alpha_2 \beta_2}{\beta_2 + 1}})$$

by setting $\lambda = O(n^{-\frac{\alpha_2}{\beta_2 + 1}})$. □

B.2. Proof for Adaptive Features

Here, we derive the loss function $\ell(\theta)$ for adaptive feature. Let $\Phi = [\phi(z_1), \dots, \phi(z_n)]$ be the feature map for Z . Then the loss $\hat{\ell}^{(Q,\lambda)}$ can be written as

$$\begin{aligned} \hat{\ell}^{(Q,\lambda)}(E, \theta) &= \text{tr} \left((\Phi - E^* \Psi_\theta) D (\Phi - E^* \Psi_\theta)^\top \right) + \lambda \|E\|^2 \\ &= \text{tr} (\Phi D \Phi^\top - 2E^* \Psi_\theta D \Phi^\top + E^* \Psi_\theta D \Psi_\theta^\top E) + \lambda \|E\|^2. \end{aligned}$$

Since for fixed θ , the minimizer of $\hat{\ell}^{(Q,\lambda)}$ with respect to E can be written as

$$\hat{E}_{Q,\lambda} = (\Psi_\theta D \Psi_\theta^\top + \lambda I)^{-1} \Psi_\theta D \Phi^\top,$$

we have

$$\begin{aligned} \hat{\ell}^{(Q,\lambda)}(\hat{E}_{Q,\lambda}, \theta) &= \text{tr} (\Phi D \Phi^\top) - 2 \text{tr} (\Phi D \Psi_\theta^\top (\Psi_\theta D \Psi_\theta^\top + \lambda I)^{-1} \Psi_\theta D \Phi^\top) \\ &\quad + \text{tr} (\Phi D \Psi_\theta^\top (\Psi_\theta D \Psi_\theta^\top + \lambda I)^{-1} \Psi_\theta D \Psi_\theta^\top (\Psi_\theta D \Psi_\theta^\top + \lambda I)^{-1} \Psi_\theta D \Phi^\top) \\ &\quad + \text{tr} (\lambda \Phi^\top D \Psi_\theta (\Psi_\theta^\top D \Psi_\theta + \lambda I)^{-1} (\Psi_\theta^\top D \Psi_\theta + \lambda I)^{-1} \Psi_\theta^\top D \Phi) \\ &= \text{tr} (\Phi D \Phi^\top) - \text{tr} (\Phi D \Psi_\theta^\top (\Psi_\theta D \Psi_\theta^\top + \lambda I)^{-1} \Psi_\theta D \Phi^\top). \end{aligned}$$

Using $G_Z = \Phi^\top \Phi$, we have

$$\begin{aligned} \ell(\theta) &= \hat{\ell}^{(Q,\lambda)}(\hat{E}_{Q,\lambda}, \theta) \\ &= \text{tr}(GD) - \text{tr} (\Phi D \Psi_\theta^\top (GD \Psi_\theta^\top + \lambda I)^{-1} \Psi_\theta D \Phi^\top) \\ &= \text{tr}(GD - GD \Psi_\theta^\top (\Psi_\theta D \Psi_\theta^\top + \lambda I)^{-1} \Psi_\theta D). \end{aligned}$$