
Conditional mean embeddings as regressors

Steffen Grünewälder¹

Guy Lever¹

Luca Baldassarre

Sam Patterson*

Arthur Gretton*[†]

Massimiliano Pontil

STEFFEN@CS.UCL.AC.UK

G.LEVER@CS.UCL.AC.UK

L.BALDASSARRE@CS.UCL.AC.UK

SAM.X.PATTERSON@GMAIL.COM

ARTHUR.GRETTON@GMAIL.COM

M.PONTIL@CS.UCL.AC.UK

CSML and *Gatsby Unit, University College London, UK, [†] MPI for Intelligent Systems

Abstract

We demonstrate an equivalence between reproducing kernel Hilbert space (RKHS) embeddings of conditional distributions and vector-valued regressors. This connection introduces a natural regularized loss function which the RKHS embeddings minimise, providing an intuitive understanding of the embeddings and a justification for their use. Furthermore, the equivalence allows the application of vector-valued regression methods and results to the problem of learning conditional distributions. Using this link we derive a sparse version of the embedding by considering alternative formulations. Further, by applying convergence results for vector-valued regression to the embedding problem we derive minimax convergence rates which are $O(\log(n)/n)$ – compared to current state of the art rates of $O(n^{-1/4})$ – and are valid under milder and more intuitive assumptions. These minimax upper rates coincide with lower rates up to a logarithmic factor, showing that the embedding method achieves nearly optimal rates. We study our sparse embedding algorithm in a reinforcement learning task where the algorithm shows significant improvement in sparsity over an incomplete Cholesky decomposition.

1. Introduction/Motivation

In recent years a framework for embedding probability distributions into reproducing kernel Hilbert spaces (RKHS)

¹**Equal contribution.** – Supplementary on arXiv.

has become increasingly popular (Smola et al., 2007). One example of this theme has been the representation of conditional expectation operators as RKHS functions, known as *conditional mean embeddings* (Song et al., 2009). Conditional expectations appear naturally in many machine learning tasks, and the RKHS representation of such expectations has two important advantages: first, conditional mean embeddings do not require solving difficult intermediate problems such as density estimation and numerical integration; and second, these embeddings may be used to compute conditional expectations directly on the basis of observed samples. Conditional mean embeddings have been successfully applied to inference in graphical models, reinforcement learning, subspace selection, and conditional independence testing (Fukumizu et al., 2008; 2009; Song et al., 2009; 2010; Grünewälder et al., 2012).

The main motivation for conditional means in Hilbert spaces has been to generalize the notion of conditional expectations from finite cases (multivariate Gaussians, conditional probability tables, and so on). Results have been established for the convergence of these embeddings in RKHS norm (Song et al., 2009; 2010), which show that conditional mean embeddings behave in the way we would hope (i.e., they may be used in obtaining conditional expectations as inner products in feature space, and these estimates are consistent under smoothness conditions). Despite these valuable results, the characterization of conditional mean embeddings remains incomplete, since these embeddings have not been defined in terms of the optimizer of a given *loss function*. This makes it difficult to extend these results, and has hindered the use of standard techniques like cross-validation for parameter estimation.

In this paper, we demonstrate that the conditional mean embedding is the solution of a vector-valued regression problem with a natural loss, resembling the standard Tikhonov regularized least-squares problem in multiple dimensions. Through this link, it is possible to access the rich the-

ory of vector-valued regression (Micchelli & Pontil, 2005; Carmeli et al., 2006; Caponnetto & De Vito, 2007; Caponnetto et al., 2008). We demonstrate the utility of this connection by providing novel characterizations of conditional mean embeddings, with important theoretical and practical implications. On the theoretical side, we establish novel convergence results for RKHS embeddings, giving a significant improvement over the rate of $O(n^{-1/4})$ due to Song et al. (2009; 2010). We derive a faster $O(\log(n)/n)$ rate which holds over large classes of probability measures, and requires milder and more intuitive assumptions. We also show our rates are optimal up to a $\log(n)$ term, following the analysis of Caponnetto & De Vito (2007). On the practical side, we derive an alternative sparse version of the embeddings which resembles the Lasso method, and provide a cross-validation scheme for parameter selection.

2. Background

In this section, we recall some background results concerning RKHS embeddings and vector-valued RKHS. For an introduction to scalar-valued RKHS we refer the reader to (Berlinet & Thomas-Agnan, 2004).

2.1. Conditional mean embeddings

Given sets \mathcal{X} and \mathcal{Y} , with a distribution P over random variables (X, Y) from $\mathcal{X} \times \mathcal{Y}$ we consider the problem of learning expectation operators corresponding to the conditional distributions $P(Y|X = x)$ on \mathcal{Y} after conditioning on $x \in \mathcal{X}$. Specifically, we begin with a kernel $L : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, with corresponding RKHS $\mathcal{H}_L \subseteq \mathbb{R}^{\mathcal{Y}}$, and study the problem of learning, for every $x \in \mathcal{X}$, the *conditional expectation mapping* $\mathcal{H}_L \ni h \mapsto \mathbb{E}[h(Y)|X = x]$. Each such map can be represented as

$$\mathbb{E}[h(Y)|X = x] = \langle h, \mu(x) \rangle_L,$$

where the element $\mu(x) \in \mathcal{H}_L$ is called the *conditional mean embedding* of $P(Y|X = x)$. Note that, for every x , $\mu(x)$ is a function on \mathcal{Y} . It is thus apparent that μ is a mapping from \mathcal{X} to \mathcal{H}_L , a point which we will expand upon shortly.

We are interested in the problem of estimating the embeddings $\mu(x)$ given an i.i.d. sample $\{(x_i, y_i)\}_{i=1}^n$ drawn from P^n . Following (Song et al., 2009; 2010), we define a second kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ with associated RKHS \mathcal{H}_K , and consider the estimate

$$\hat{\mu}(x) := \sum_{i=1}^n \alpha_i(x) L(y_i, \cdot), \quad (1)$$

where $\alpha_i(x) = \sum_{j=1}^n W_{ij} K(x_j, x)$, and where $\mathbf{W} := (\mathbf{K} + \lambda n \mathbf{I})^{-1}$, $\mathbf{K} = (K(x_i, x_j))_{i,j=1}^n$, and λ is a chosen regularization parameter. This expression suggests that

the conditional mean embedding is the solution to an underlying regression problem: we will formalize this link in Section 3. In the remainder of the present section, we introduce the necessary terminology and theory for vector valued regression in RHKSs.

2.2. Vector-valued regression and RKHSs

We recall some background on learning vector-valued functions using kernel methods (see Micchelli & Pontil, 2005, for more detail). We are given a sample $\{(x_i, v_i)\}_{i \leq m}$ drawn i.i.d. from some distribution over $\mathcal{X} \times \mathcal{V}$, where \mathcal{X} is a non-empty set and $(\mathcal{V}, \langle \cdot, \cdot \rangle_{\mathcal{V}})$ is a Hilbert space. Our goal is to find a function $f : \mathcal{X} \rightarrow \mathcal{V}$ with low error, as measured by

$$\mathbb{E}_{(X, V)} [\|f(X) - V\|_{\mathcal{V}}^2]. \quad (2)$$

This is the *vector-valued regression* problem (square loss).

One approach to the vector-valued regression problem is to model the regression function as being in a vector-valued RKHS of functions taking values in \mathcal{V} , which can be defined by analogy with the scalar valued case.

Definition A Hilbert space $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\Gamma})$ of functions $h : \mathcal{X} \rightarrow \mathcal{V}$ is an RKHS if for all $x \in \mathcal{X}, v \in \mathcal{V}$ the linear functional $h \mapsto \langle v, h(x) \rangle_{\mathcal{V}}$ is continuous.

The reproducing property for vector-valued RHKSs follows from this definition (see Micchelli & Pontil, 2005, Sect. 2). By the Riesz representation theorem, for each $x \in \mathcal{X}$ and $v \in \mathcal{V}$, there exists a linear operator from \mathcal{V} to \mathcal{H}_{Γ} written $\Gamma_x v \in \mathcal{H}_{\Gamma}$, such that for all $h \in \mathcal{H}_{\Gamma}$,

$$\langle v, h(x) \rangle_{\mathcal{V}} = \langle h, \Gamma_x v \rangle_{\Gamma}.$$

It is instructive to compare to the scalar-valued RKHS \mathcal{H}_K , for which the linear operator of evaluation δ_x mapping $h \in \mathcal{H}_K$ to $h(x) \in \mathbb{R}$ is continuous: then Riesz implies there exists a K_x such that $yh(x) = \langle h, yK_x \rangle_K$.

We next introduce the vector-valued reproducing kernel, and show its relation to Γ_x . Writing as $\mathcal{L}(\mathcal{V})$ the space of bounded linear operators from \mathcal{V} to \mathcal{V} , the reproducing kernel $\Gamma(x, x') \in \mathcal{L}(\mathcal{V})$ is defined as

$$\Gamma(x, x')v = (\Gamma_{x'}v)(x) \in \mathcal{V}.$$

From this definition and the reproducing property, the following holds (Micchelli & Pontil, 2005, Prop. 2.1).

Proposition 2.1. *A function $\Gamma : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{L}(\mathcal{V})$ is a kernel if it satisfies: (i) $\Gamma(x, x') = \Gamma(x', x)^*$, (ii) for all $n \in \mathbb{N}, \{(x_i, v_i)\}_{i \leq n} \subseteq \mathcal{X} \times \mathcal{V}$ we have that $\sum_{i, j \leq n} \langle v_i, \Gamma(x_i, x_j)v_j \rangle_{\mathcal{V}} \geq 0$.*

It is again helpful to consider the scalar case: here, $\langle K_x, K_{x'} \rangle_K = K(x, x')$, and to every positive definite

kernel $K(x, x')$ there corresponds a unique (up to isometry) RKHS for which K is the reproducing kernel. Similarly, if $\Gamma : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{L}(\mathcal{V})$ is a kernel in the sense of Proposition 2.1, there exists a unique (up to isometry) RKHS, with Γ as its reproducing kernel (Micchelli & Pontil, 2005, Th. 2.1). Furthermore, the RKHS \mathcal{H}_Γ can be described as the RKHS limit of finite sums; that is, \mathcal{H}_Γ is up to isometry equal to the closure of the linear span of the set $\{\Gamma_x v : x \in \mathcal{X}, v \in \mathcal{V}\}$, wrt the RKHS norm $\|\cdot\|_\Gamma$.

Importantly, it is possible to perform regression in this setting. One approach to the vector-valued regression problem is to replace the unknown true error (2) with a sample-based estimate $\sum_{i=1}^n \|v_i - f(x_i)\|_{\mathcal{V}}^2$, restricting f to be an element of an RKHS \mathcal{H}_Γ (of vector-valued functions), and regularizing w.r.t. the \mathcal{H}_Γ norm, to prevent overfitting. We thus arrive at the following regularized empirical risk,

$$\widehat{\mathcal{E}}_\lambda(f) := \sum_{i=1}^n \|v_i - f(x_i)\|_{\mathcal{V}}^2 + \lambda \|f\|_\Gamma^2. \quad (3)$$

Theorem 2.2. (Micchelli & Pontil, 2005, Th. 4) *If f^* minimises $\widehat{\mathcal{E}}_\lambda$ in \mathcal{H}_Γ then it is unique and has the form,*

$$f^* = \sum_{i=1}^n \Gamma_{x_i} c_i,$$

where the coefficients $\{c_i\}_{i \leq n}$, $c_i \in \mathcal{V}$ are the unique solution of the system of linear equations

$$\sum_{i \leq n} (\Gamma(x_j, x_i) + \lambda \delta_{ji}) c_i = v_j, \quad 1 \leq j \leq n.$$

In the scalar case we have that $|f(x)| \leq \sqrt{K(x, x)} \|f\|_K$. Similarly it holds that $\|f(x)\| \leq \|\Gamma(x, x)\| \|f\|_\Gamma$, where $\|\cdot\|$ denotes the operator norm (Micchelli & Pontil, 2005, Prop. 1). Hence, if $\|\Gamma(x, x)\| \leq B$ for all x then

$$\|f(x)\|_{\mathcal{V}} \leq B \|f\|_\Gamma. \quad (4)$$

Finally, we need a result that tells us when all functions in an RKHS are continuous. In the scalar case this is guaranteed if $K(x, \cdot)$ is continuous for all x and K is bounded. In our case we have (Carmeli et al., 2006)[Prop. 12]:

Corollary 2.3. *If \mathcal{X} is a Polish space, \mathcal{V} a separable Hilbert space and the mapping $x \mapsto \Gamma(\cdot, x)$ is continuous, then \mathcal{H}_Γ is a subset of the set of continuous functions from \mathcal{X} to \mathcal{V} .*

3. Estimating conditional expectations

In this section, we show the problem of learning conditional mean embeddings can be naturally formalised in the framework of vector-valued regression, and in doing so we derive an equivalence between the conditional mean embeddings and a vector-valued regressor.

3.1. The equivalence between conditional mean embeddings and a vector-valued regressor

Conditional expectations $\mathbb{E}[h(Y)|X = x]$ are linear in the argument h so that, when we consider $h \in \mathcal{H}_L$, the Riesz representation theorem implies the existence of an element $\mu(x) \in \mathcal{H}_L$ such that $\mathbb{E}[h(Y)|X = x] = \langle h, \mu(x) \rangle_L$ for all h . That being said, the dependence of μ on x may be complicated. A natural optimisation problem associated to this approximation problem is to therefore find a function $\mu : \mathcal{X} \rightarrow \mathcal{H}_L$ such that the following objective is small

$$\mathcal{E}[\mu] := \sup_{\|h\|_L \leq 1} \mathbb{E}_X \left[(\mathbb{E}_Y[h(Y)|X] - \langle h, \mu(X) \rangle_L)^2 \right]. \quad (5)$$

Note that the risk function cannot be used directly for estimation, because we do not observe $\mathbb{E}_Y[h(Y)|X]$, but rather pairs (X, Y) drawn from P . However, we can bound this risk function with a surrogate risk function that has a sample based version,

$$\begin{aligned} & \sup_{\|h\|_L \leq 1} \mathbb{E}_X \left[(\mathbb{E}_Y[h(Y)|X] - \langle h, \mu(X) \rangle_L)^2 \right] \\ &= \sup_{\|h\|_L \leq 1} \mathbb{E}_X \left[(\mathbb{E}_Y[\langle h, L(Y, \cdot) \rangle_L | X] - \langle h, \mu(X) \rangle_L)^2 \right] \\ &\leq \sup_{\|h\|_L \leq 1} \mathbb{E}_{X, Y} [\langle h, L(Y, \cdot) - \mu(X) \rangle_L^2] \\ &\leq \sup_{\|h\|_L \leq 1} \|h\|_L^2 \mathbb{E}_{X, Y} [\|L(Y, \cdot) - \mu(X)\|_L^2] \\ &= \mathbb{E}_{(X, Y)} [\|L(Y, \cdot) - \mu(X)\|_L^2], \end{aligned} \quad (6)$$

where the first and second bounds follow by Jensen's and Cauchy-Schwarz's inequalities, respectively. Let us denote this surrogate risk function as

$$\mathcal{E}_s[\mu] := \mathbb{E}_{(X, Y)} [\|L(Y, \cdot) - \mu(X)\|_L^2]. \quad (7)$$

The two risk functions \mathcal{E} and \mathcal{E}_s are closely related and in Section 3.3 we examine their relation.

We now replace the expectation in (6) with an empirical estimate, to obtain the sample-based loss,

$$\widehat{\mathcal{E}}_n[\mu] := \sum_{i=1}^n \|L(y_i, \cdot) - \mu(x_i)\|_L^2. \quad (8)$$

Taking (8) as our empirical loss, then following Section 2.2 we add a regularization term to provide a well-posed problem and prevent overfitting,

$$\widehat{\mathcal{E}}_{\lambda, n}[\mu] := \sum_{i=1}^n \|L(y_i, \cdot) - \mu(x_i)\|_L^2 + \lambda \|\mu\|_\Gamma^2. \quad (9)$$

We denote the minimizer of (9) by $\hat{\mu}_{\lambda, n}$,

$$\hat{\mu}_{\lambda, n} := \underset{\mu}{\operatorname{argmin}} \left\{ \widehat{\mathcal{E}}_{\lambda, n}[\mu] \right\}. \quad (10)$$

Thus, recalling (3), we can see that the problem (10) is posed as a vector-valued regression problem with the training data now considered as $\{(x_i, L(y_i, \cdot))\}_{i=1}^n$ (and we identify \mathcal{H}_L with the general Hilbert space \mathcal{V} of Section 2.2). From Theorem 2.2, the solution is

$$\hat{\mu}_{\lambda, n} = \sum_{i=1}^n \Gamma_{x_i} c_i, \quad (11)$$

where the coefficients $\{c_i\}_{i \leq n}$, $c_i \in \mathcal{H}_L$ are the unique solution of the linear equations

$$\sum_{i \leq n} (\Gamma(x_j, x_i) + \lambda \delta_{ji}) c_i = L(y_j, \cdot), \quad 1 \leq j \leq n.$$

It remains to choose the kernel Γ . Given a real-valued kernel K on \mathcal{X} , a natural choice for the RKHS \mathcal{H}_Γ would be the space of functions from \mathcal{X} to \mathcal{H}_L whose elements are defined as functions via $(h, K(x, \cdot))(x') := K(x, x')h$, which is isomorphic to $\mathcal{H}_L \otimes \mathcal{H}_K$, with inner product

$$\langle gK(x, \cdot), hK(x', \cdot) \rangle_\Gamma := \langle g, h \rangle_L K(x, x') \quad (12)$$

for all $g, h \in \mathcal{H}_L$. Its easy to check that this satisfies the conditions to be a vector-valued RKHS— in fact it corresponds to the choice $\Gamma(x, x') = K(x, x')\text{Id}$, where $\text{Id} : \mathcal{H}_L \rightarrow \mathcal{H}_L$ is the identity map on \mathcal{H}_L . The solution to (10) with this choice is then given by (11), with

$$\begin{aligned} \sum_{i \leq n} (K(x_j, x_i) + \lambda \delta_{ji}) c_i &= L(y_j, \cdot), \quad 1 \leq j \leq n \\ c_i &= \sum_{j \leq n} W_{ij} L(y_j, \cdot), \quad 1 \leq i \leq n, \end{aligned}$$

where $\mathbf{W} = (\mathbf{K} + \lambda \mathbf{I})^{-1}$, which corresponds exactly the embeddings (1) presented in (Song et al., 2009; 2010) (after a rescaling of λ). Thus we have shown that the embeddings of Song et al. are the solution to a regression problem for a particular choice of operator-valued kernel. Further, the loss defined by (7) is a key error functional in this context since it is the objective which the estimated embeddings attempt to minimise. In Sec. 3.3 we will see that this does not always coincide with (5) which may be a more natural choice. In Sec. 4 we analyze the performance of the embeddings defined by (10) at minimizing the objective (7).

3.2. Some consequences of this equivalence

We derive some immediate benefits from the connection described above. Since the embedding problem has been identified as a regression problem with training set $\mathcal{D} := \{(x_i, L(y_i, \cdot))\}_{i=1}^m$, we can define a cross validation scheme for parameter selection in the usual way: by holding out a subsample $\mathcal{D}_{\text{val}} = \{(x_{t_j}, L(y_{t_j}, \cdot))\}_{j=1}^T \subset \mathcal{D}$, we can train embeddings $\hat{\mu}$ on $\mathcal{D} \setminus \mathcal{D}_{\text{val}}$ over a grid of kernel or

Input/Output space	(i) \mathcal{X} is Polish. (ii) \mathcal{V} is separable. (f.b.a.) (iii) $\exists C > 0$ such that $\forall x \in \mathcal{X}$ $\text{Tr}(\Gamma(x, x)) \leq C$ holds.
Space of regressors	(iv) \mathcal{H}_Γ is separable. (v) All Γ_x^* are HS. (f.b.a.) (vi) $B : (x, y) \rightarrow \langle f, \Gamma(y, x)g \rangle_{\mathcal{V}}$ is measurable $\forall f, g \in \mathcal{V}$.
True distribution	(vii) $L(y, y) < \infty$ for all $y \in \mathcal{Y}$. (viii) $\mathcal{E}_s[\mu^*] = \inf_{\mu \in \mathcal{H}_\Gamma} \mathcal{E}_s[\mu]$.

Table 1. Assumptions for Corollary 4.1 and 4.2. f.b.a. stands for fulfilled by assumption that \mathcal{V} is finite dimensional.

regularization parameters, choosing the parameters achieving the best error $\sum_{j=1}^T \|\hat{\mu}(x_{t_j}) - L(y_{t_j}, \cdot)\|_L^2$ on the validation set (or over many folds of cross validation). Another key benefit will be a much improved performance analysis of the embeddings, presented in Section 4.

3.3. Relations between the error functionals \mathcal{E} and \mathcal{E}_s

In Section 3.1 we introduced an alternative risk function \mathcal{E}_s for \mathcal{E} , which we used to derive an estimation scheme to recover conditional mean embeddings. We now examine the relationship between the two risk functionals. When the true conditional expectation on functions $h \in \mathcal{H}_L$ can be represented through an element $\mu^* \in \mathcal{H}_\Gamma$ then μ^* minimises both objectives.

Theorem 3.1 (Proof in App. A). *If there exists a $\mu^* \in \mathcal{H}_\Gamma$ such that for any $h \in \mathcal{H}_L$: $\mathbb{E}[h|X] = \langle h, \mu^*(X) \rangle_L$ $P_{\mathcal{X}}$ -a.s., then μ^* is the $P_{\mathcal{X}}$ -a.s. unique minimiser of both objectives:*

$$\mu^* = \underset{\mu \in \mathcal{H}_\Gamma}{\text{argmin}} \mathcal{E}[\mu] = \underset{\mu \in \mathcal{H}_\Gamma}{\text{argmin}} \mathcal{E}_s[\mu] \quad P_{\mathcal{X}} \text{ a.s.}$$

Thus in this case, the embeddings of Song et al. (e.g. 2009) minimise both (5) and (7). More generally, however, this may not be the case. Let us define an element $\hat{\mu}$ that is δ close w.r.t. the error \mathcal{E}_s to the minimizer μ' of \mathcal{E}_s in \mathcal{H}_Γ (this might for instance be the minimizer of the empirical regularized loss for sufficiently many samples). We are interested in finding conditions under which $\mathcal{E}(\hat{\mu})$ is not much worse than a good approximation μ^* in \mathcal{H}_Γ to the conditional expectation. The sense in which μ^* approximates the conditional expectation is somewhat subtle: μ^* must closely approximate the conditional expectation of functions $\mu \in \mathcal{H}_\Gamma$ under the original loss \mathcal{E} (note that the loss \mathcal{E} was originally defined in terms of functions $h \in \mathcal{H}_L$).

Theorem 3.2 (Proof in App. A). *Let μ' be a minimiser of \mathcal{E}_s and $\hat{\mu}$ be an element of \mathcal{H}_Γ with $\mathcal{E}_s[\hat{\mu}] \leq \mathcal{E}_s[\mu'] + \delta$. Define, $\mathcal{A} := \{(\eta, \hat{\mu}) \mid \eta^2 =$*

$\sup_{\|\mu\|_{\Gamma} \leq 1} \mathbb{E}_X [\mathbb{E}_Y[\mu(X)|X] - \langle \mu(X), \tilde{\mu}(X) \rangle_L]^2$, then

$$\mathcal{E}[\hat{\mu}] \leq \inf_{(\eta, \mu^*) \in \mathcal{A}} \left(\sqrt{\mathcal{E}[\mu^*]} + \sqrt{8\eta(\|\mu^*\|_{\Gamma} + \|\hat{\mu}\|_{\Gamma})} + \delta^{\frac{1}{2}} \right)^2.$$

Apart from the more obvious condition that δ be small, the above theorem suggests that $\|\hat{\mu}\|_{\Gamma}$ should also be made small for the solution $\hat{\mu}$ to have low error \mathcal{E} . In other words, even in the infinite sample case, the regularization of $\hat{\mu}$ in \mathcal{H}_{Γ} is important.

4. Better convergence rates for embeddings

The interpretation of the mean embedding as a vector valued regression problem allows us to apply regression minimax theorems to study convergence rates of the embedding estimator. These rates are considerably better than the current state of the art for the embeddings, and hold under milder and more intuitive assumptions.

We start by comparing the statements which we derive from (Caponnetto & De Vito, 2007, Thm.s 1 and 2) with the known convergence results for the embedding estimator. We follow this up with a discussion of the rates and a comparison of the assumptions.

4.1. Convergence theorems

We address the performance of the embeddings defined by (10) in terms of asymptotic guarantees on the loss \mathcal{E}_s defined by (7). Caponnetto & De Vito (2007) study uniform convergence rates for regression. Convergence rates of learning algorithms can not be uniform on the set of all probability distributions if the output vector space is an infinite dimensional RKHS (Caponnetto & De Vito, 2007)[p. 4]. It is therefore necessary to restrict ourselves to a subset of probability measures. This is done by Caponnetto & De Vito (2007) by defining families of probability measures $\mathcal{P}(b, c)$ indexed by two parameters $b \in [1, \infty]$ and $c \in [1, 2]$. We discuss the family $\mathcal{P}(b, c)$ in detail below. The important point at the moment is that b and c affect the optimal schedule for the regulariser λ and the convergence rate. The rate of convergence is better for higher b and c values. Caponnetto & De Vito (2007) provide convergence rates for all choices of b and c . We restrict ourself to the best case $b = \infty, c > 1$ and the worst case¹ $b = 2, c = 1$.

We recall that the estimated conditional mean embeddings $\hat{\mu}_{\lambda, n}$ are given by (10), where λ is a chosen regularization parameter. We assume λ_n is chosen to follow a specific schedule, dependent upon n : we denote by $\hat{\mu}_n$ the embeddings following this schedule and $\mu' := \operatorname{argmin}_{\mu \in \mathcal{H}_{\Gamma}} \mathcal{E}_s[\mu]$. Given this optimal rate of decrease for λ_n , Thm. 1 of Caponnetto & De Vito (2007) yields the

¹Strictly speaking the worst case is $b \downarrow 1$ (see supp.).

following convergence statements for the estimated embeddings, under assumptions to be discussed in Section 4.2.

Corollary 4.1. *Let $b = \infty, c > 1$ then for every $\epsilon > 0$ there exists a constant τ such that*

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}(b, c)} P^n \left[\mathcal{E}_s[\hat{\mu}_n] - \mathcal{E}_s[\mu'] > \tau \frac{1}{n} \right] < \epsilon.$$

Let $b = 2$ and $c = 1$ then for every $\epsilon > 0$ there exists a constant τ such that

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}(b, c)} P^n \left[\mathcal{E}_s[\hat{\mu}_n] - \mathcal{E}_s[\mu'] > \tau \left(\frac{\log n}{n} \right)^{\frac{2}{3}} \right] < \epsilon.$$

The rate for the estimate $\hat{\mu}_n$ can be complemented with minimax lower rates for vector valued regression (Caponnetto & De Vito, 2007)[Th. 2] in the case that $b < \infty$.

Corollary 4.2. *Let $b = 2$ and $c = 1$ and let $\Lambda_n := \{l_n | l_n : (\mathcal{X} \times \mathcal{Y})^n \rightarrow \mathcal{H}_{\Gamma}\}$ be the set of all learning algorithm working on n samples, outputting $\nu_n \in \mathcal{H}_{\Gamma}$. Then for every $\epsilon > 0$ there exists a constant $\tau > 0$ such that*

$$\liminf_{n \rightarrow \infty} \inf_{l_n \in \Lambda_n} \sup_{P \in \mathcal{P}(b, c)} P^n \left[\mathcal{E}_s[\nu_n] - \mathcal{E}_s[\mu'] > \tau \left(\frac{1}{n} \right)^{\frac{2}{3}} \right] > 1 - \epsilon.$$

This corollary tells us that there exists no learning algorithm which can achieve better rates than $n^{-\frac{2}{3}}$ uniformly over $\mathcal{P}(2, 1)$, and hence the estimate $\hat{\mu}_n$ is optimal up to a logarithmic factor.

State of the art results for the embedding The current convergence result for the embedding is proved by Song et al. (2010, Th.1). A crucial assumption that we discuss in detail below is that the mapping $x \mapsto \mathbb{E}[h(Y)|X = x]$ is in the RKHS \mathcal{H}_K of the real valued kernel, i.e. that for all $h \in \mathcal{H}_L$ we have that there exists a $f_h \in \mathcal{H}_K$, such that

$$\mathbb{E}[h(Y)|X = x] = f_h(x). \quad (13)$$

The result of Song et al. implies the following (see App. C): if $K(x, x) < B$ for all $x \in \mathcal{X}$ and the schedule $\lambda(n) = n^{-1/4}$ is used: for a fixed probability measure P , there exists a constant τ such that

$$\lim_{n \rightarrow \infty} P^n \left[\mathcal{E}[\hat{\mu}_n] > \tau \left(\frac{1}{n} \right)^{\frac{1}{4}} \right] = 0, \quad (14)$$

where $\hat{\mu}_n(x)$ is the estimate from Song et al. No complementary lower bounds were known until now.

Comparison The first thing to note is that under the assumption that $\mathbb{E}[h|X]$ is in the RKHS \mathcal{H}_K the minimiser of \mathcal{E}_s and \mathcal{E} are a.e. equivalent due to Theorem 3.1: the assumption implies a $\mu^* \in \mathcal{H}_\Gamma$ exists with $\mathbb{E}[h|X] = \langle h, \mu^*(X) \rangle_L$ for all $h \in \mathcal{V}$ (see App. B.4 for details). Hence, under this assumption, the statements from eq. 14 and Cor. 4.1 ensure we converge to the true conditional expectation, and achieve an error of 0 in the risk \mathcal{E} .

In the case that this assumption is not fulfilled and eq. 14 is not applicable, Cor. 4.1 still tells us that we converge to the minimiser of \mathcal{E}_s . Coupling this statement with Thm. 3.2 allows us to bound the distance to the minimal error $\mathcal{E}[\mu^*]$, where $\mu^* \in \mathcal{H}_\Gamma$ minimises \mathcal{E} .

The other main differences are obviously the rates, and that Cor. 4.1 bounds the error uniformly over a space of probability measures, while eq. 14 provides only a point-wise statement (i.e., for a fixed probability measure P).

4.2. Assumptions

Cor. 4.1 and 4.2 Our main assumption is that \mathcal{H}_L is finite dimensional. It is likely that this assumption can be weakened, but this requires a deeper analysis.

The assumptions of Caponnetto & De Vito (2007) are summarized in Table 1, where we provide details in App. B.2. App. B.1 contains simple and complete assumptions that ensure all statements in the table hold. Beside some measure theoretic issues, the assumptions are fulfilled if for example, 1) \mathcal{X} is a compact subset of \mathbb{R}^n , \mathcal{Y} is compact, \mathcal{H}_L is a finite dimensional RKHS, Γ and L are continuous; 2) $\mathcal{E}_s[\mu'] = \inf_{\mu \in \mathcal{H}_\Gamma} \mathcal{E}_s[\mu]$. This last condition is unintuitive, but can be rewritten in the following way:

Theorem 4.3 (Proof in App.). *Let $\|h\|_{\mathcal{V}}, \|\mu(x) - h\|_{\mathcal{V}}$ be integrable for all $h \in \mathcal{H}_\Gamma$ and let \mathcal{V} be finite dimensional. Then there exists a $\mu' \in \mathcal{H}_\Gamma$ with $\mathcal{E}_s[\mu'] = \inf_{\mu \in \mathcal{H}_\Gamma} \mathcal{E}_s[\mu]$ iff a $B > 0$ exists and a sequence $\{\mu_n\}_{n \in \mathbb{N}}$ with $\mathcal{E}_s[\mu_n] \leq \inf_{\mu \in \mathcal{H}_\Gamma} \mathcal{E}_s[\mu] + 1/n$ and $\|\mu_n\|_\Gamma < B$.*

The intuition is that the condition is not fulfilled if we need to make μ_n more and more complex (in the sense of a high RKHS norm) to optimize the risk.

Definition and discussion of $\mathcal{P}(b, c)$ The family of probability measures $\mathcal{P}(b, c)$ is characterized through spectral properties of the kernel function Γ . The assumptions correspond to assumptions on the eigenvalues in Mercer’s theorem in the real valued case, i.e. that there are finitely many eigenvalues or that the eigenvalues decrease with a certain rate. In detail, define the operator A through $A(\phi)(x') := \int_{\mathcal{X}} \Gamma(x', x)\phi(x)dP_{\mathcal{X}}$, where $\phi \in L^2(P_{\mathcal{X}})$. A can be written as (Caponnetto & De Vito, 2007, Rem. 2) $A = \sum_{n=1}^N \lambda_n \langle \cdot, \phi_n \rangle_P \phi_n$, where the inner product is the L^2 inner product with measure $P_{\mathcal{X}}$ and $N = \infty$ is al-

lowed. As in the real valued case, the eigendecomposition depends on the measure on the space \mathcal{X} but is independent of the distribution on \mathcal{Y} . The eigendecomposition measures the complexity of the kernel, where the lowest complexity is achieved for finite N — that is, the case $b = \infty, c > 1$ — and has highest complexity if the eigenvalues decrease with the slowest possible rate, $\lambda_n < C/n$ for a constant C . The case $b = 2, c = 1$ correspond to a slightly faster decay, namely, $\lambda_n < C/n^2$. In essence, there are no assumptions on the distribution on \mathcal{Y} , but only on the complexity of the kernel Γ as measured with $P_{\mathcal{X}}$.

Embedding The results of Song et al. (2010) do not rely on the assumption that \mathcal{V} is finite dimensional. Other conditions on the distribution are required, however, which are challenging to verify. To describe these conditions, we recall the real-valued RKHS \mathcal{H}_K with kernel K , and define the uncentred cross-covariance operator $C_{YX} : \mathcal{H}_K \rightarrow \mathcal{H}_L$ such that $\langle g, C_{YX}f \rangle_{\mathcal{H}_L} = \mathbb{E}_{XY}(f(X)g(Y))$, with the covariance operator C_{XX} defined by analogy. One of the two main assumptions of Song et al. is that $C_{YX}C_{XX}^{-3/2}$ needs to be Hilbert-Schmidt. The covariances C_{YX} and C_{XX} are compact operators, meaning C_{XX} is not invertible when \mathcal{H}_K is infinite dimensional (this gives rise to a notational issue, although the “product” operator $C_{YX}C_{XX}^{-3/2}$ may still be defined). Whether $C_{YX}C_{XX}^{-3/2}$ is Hilbert-Schmidt (or even bounded) will depend on the underlying distribution P_{XY} and on the kernels K and L . At this point, however, there is no easy way to translate properties of P_{XY} to guarantees that the assumption holds.

The second main assumption is that the conditional expectation can be represented as an RKHS element (see App B.4). Even for rich RKHSs (such as universal RKHSs), it can be challenging to determine the associated conditions on the distribution P_{XY} . For simple finite dimensional RKHSs, the assumption may fail, as shown below.

Corollary 4.4 (Proof in App. C). *Let \mathcal{V} be finite dimensional such that a function $h \in \mathcal{V}$ exists with $h(y) \geq \epsilon > 0$ for all $y \in \mathcal{Y}$. Furthermore, let $\mathcal{X} := [-1, 1]$ and the reproducing kernel for \mathcal{H}_K be $K(x, y) = xy$. Then there exists no measure for which the assumption from eq. (13) can be fulfilled.*

5. Sparse embeddings

In many practical situations it is desirable to approximate the conditional mean embedding by a sparse version which involves a smaller number of parameters. For example, in the context of reinforcement learning and planning, the sample size n is large and we want to use the embeddings over and over again, possibly on many different tasks and over a long time frame.

Here we present a technique to achieve a sparse approximation of the sample mean embedding. Recall that this is given by the formula (cf. equation (11))

$$\hat{\mu}(x) = \sum_{i,j=1}^n W_{ij} K(x_i, x) L(y_j, \cdot),$$

where $W = (K + n\lambda I)^{-1}$. A natural approach to find a sparse approximation of $\hat{\mu}$ is to look for a function μ which is close to $\hat{\mu}$ according to the RKHS norm $\|\cdot\|_{\Gamma}$ (in App. D we establish a link between this objective and our cost function \mathcal{E}). In the special case that $\Gamma = \text{KId}$ this amounts to solving the optimization problem

$$\min_{M \in \mathbb{R}^{n \times n}} f(M - W) + \gamma \|M\|_{1,1} \quad (15)$$

where γ is a positive parameter, $\|W\|_{1,1} := \sum_{i,j}^n |M_{ij}|$ and

$$f(M) = \left\| \sum_{i,j=1}^n M_{ij} K(x_i, \cdot) L(y_j, \cdot) \right\|_{K \otimes L}. \quad (16)$$

Problem (15) is equivalent to a kind of Lasso problem with n^2 variables: when $\gamma = 0$, $M = W$ at the optimum and the approximation error is zero, however as γ increases, the approximation error increases as well, but the solution obtained becomes sparse (many of the elements of matrix M are equal to zero).

A direct computation yields that the above optimization problem is equivalent to

$$\min_{M \in \mathbb{R}^{n \times n}} \text{tr}((M - W)^{\top} K (M - W) L) + \gamma \sum_{i,j=1}^n |M_{ij}|. \quad (17)$$

In the experiments in Section 5.1, we solve problem (17) with FISTA (Beck & Teboulle, 2009), an optimal first order method which requires $O(1/\sqrt{\epsilon})$ iterations to reach a ϵ accuracy of the minimum value in (17), with a cost per iteration of $O(n^2)$ in our case. The algorithm is outlined below, where $S_{\gamma}(Z_{ij}) = \text{sign}(Z_{ij})(|Z_{ij}| - \gamma)_+$ and $(z)_+ = z$ if $z > 0$ and zero otherwise.

Algorithm 1 LASSO-like Algorithm

input: W, γ, K, L **output:** M
 $Z_1 = Q_1 = 0, \theta_1 = 1, C = \|K\| \|L\|$
for $t=1, 2, \dots$ **do**
 $Z_{t+1} = S_{\gamma C}(Q_t - C(K Q_t L - G W L))$
 $\theta_{t+1} = \frac{1 + \sqrt{1 + 4\theta_t^2}}{2}$
 $Q_{t+1} = Z_{t+1} + \frac{\theta_t - 1}{\theta_{t+1}}(Z_t - Z_{t+1})$
end for

Other sparsity methods could also be employed. For example, we may replace the norm $\|\cdot\|_{1,1}$ by a block ℓ_1/ℓ_2

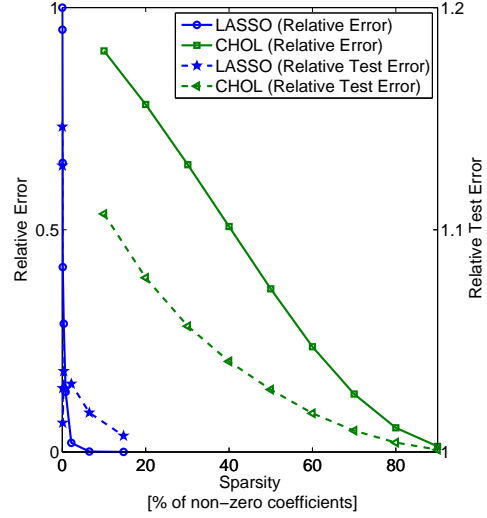


Figure 1. Comparison between our sparse algorithm and an incomplete Cholesky decomposition. The x-axis shows the level of sparsity, where on the right side the original solution is recovered. The y-axis shows the distance to the dense solution and the test error relative to the dense solution.

norm. That is, we may choose the norm $\|M\|_{2,1} := \sum_{j=1}^n \sqrt{\sum_{i=1}^n M_{ij}^2}$, which is the sum of the ℓ_2 norms of the rows of M . This penalty function encourages sparse approximations which use few input points but all the outputs. Similarly, the penalty $\|M^{\top}\|_{2,1}$ will sparsify over the outputs. Finally, if we wish to remove many pair of examples we may use the more sophisticated penalty $\sum_{i,j=1}^n \sqrt{\sum_{k=1}^m M_{ik}^2 + M_{kj}^2}$.

5.1. Experiment

We demonstrate here that the link between vector-valued regression and the mean embeddings can be leveraged to develop useful embedding alternatives that exploit properties of the regression formulation: we apply the sparse algorithm to a challenging reinforcement learning task. The sparse algorithm makes use of the labels, while other algorithms to sparsify the embeddings without our regression interpretation cannot make use of these. In particular, a popular method to sparsify is the incomplete Cholesky decomposition (Shawe-Taylor & Cristianini, 2004), which sparsifies based on the distribution on the input space \mathcal{X} only. We compare to this method in the experiments.

The reinforcement learning task is the under-actuated pendulum swing-up problem from Deisenroth et al. (2009). We generate a discrete-time approximation of the continuous-time pendulum dynamics as done in (Deisenroth et al., 2009). Starting from an arbitrary state the goal is to swing the pendulum up and balance it in the inverted position. The applied torque is $u \in [-5, 5]Nm$ and is not

sufficient for a direct swing up. The state space is defined by the angle $\theta \in [-\pi, \pi]$ and the angular velocity, $\omega \in [-7, 7]$. The reward is given by the function $r(\theta, \omega) = \exp(-\theta^2 - 0.2\omega^2)$. The learning algorithm is a kernel method which uses the mean embedding estimator to perform policy iteration (Grünwälder et al., 2012). Sparse solutions are in this task very useful as the policy iteration applies the mean embedding many times to perform updates. The input space has 4 dimensions (sine and cosine of the angle, angular velocity and applied torque), while the output has 3 (sine and cosine of the angle and angular velocity). We sample uniformly a training set of 200 examples and learn the mean conditional embedding using the direct method (1). We then compare the sparse approximation obtained by Algorithm 1 using different values of the parameter γ to the approximation obtained via an incomplete Cholesky decomposition at different levels of sparsity. We assess the approximations using the test error and (16), which is an upper bound on the generalization error (see App. D) and report the results in Figure 1.

6. Outlook

We have established a link between vector-valued regression and conditional mean embeddings. On the basis of this link, we derived a sparse embedding algorithm, showed how cross-validation can be performed, established better convergence rates under milder assumptions, and complemented these upper rates with lower rates, showing that the embedding estimator achieves near optimal rates.

There are a number of interesting questions and problems which follow from our framework. It may be valuable to employ other kernels Γ in place of the kernel $K(x, y)\text{Id}$ that leads to the mean embedding, so as to exploit knowledge about the data generating process. As a related observation, for the kernel $\Gamma(x, y) := K(x, y)\text{Id}$, Γ_x is not a Hilbert-Schmidt operator if \mathcal{V} is infinite dimensional, as

$$\|\Gamma_x\|_{HS}^2 = K(x, x) \sum_{i=1}^{\infty} \langle e_i, \text{Id } e_i \rangle_L = \infty,$$

however the convergence results from (Caponnetto & De Vito, 2007) assume Γ_x to be Hilbert-Schmidt. While this might simply be a result of the technique used in (Caponnetto & De Vito, 2007), it might also indicate a deeper problem with the standard embedding estimator, namely that if \mathcal{V} is infinite dimensional then the rates degrade. The latter case would have a natural interpretation as an overfitting effect, as Id does not “smooth” the element $h \in \mathcal{V}$.

Our sparsity approach can potentially be equipped with other regularisers that cut down on rows and columns of the W matrix in parallel. Certainly, ours is not the only sparse regression approach, and other sparse regularizers might yield good performance on appropriate problems.

Acknowledgements

The authors want to thank for the support of the EPSRC #EP/H017402/1 (CARDyAL) and the European Union #FP7-ICT-270327 (Complacs).

References

- Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Imaging Sciences*, 2(1):183–202, 2009.
- Berlinet, A. and Thomas-Agnan, C. *Reproducing kernel Hilbert spaces in Probability and Statistics*. Kluwer, 2004.
- Caponnetto, A. and De Vito, E. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- Caponnetto, A., Micchelli, C.A., Pontil, M., and Ying, Y. Universal multi-task kernels. *JMLR*, 9, 2008.
- Carmeli, C., De Vito, E., and Toigo, A. Vector valued reproducing kernel Hilbert spaces of integrable functions and mercer theorem. *Analysis and Applications*, 4(4):377–408, 2006.
- Deisenroth, M. P., Rasmussen, C. E., and Peters, J. Gaussian process dynamic programming. *Neurocomputing*, 72(7-9), 2009.
- Fremlin, D.H. *Measure Theory Volume 1: The Irreducible Minimum*. Torres Fremlin, 2000.
- Fremlin, D.H. *Measure Theory Volume 4: Topological Measure Spaces*. Torres Fremlin, 2003.
- Fukumizu, K., Gretton, A., Sun, X., and Schölkopf, B. Kernel measures of conditional dependence. In *NIPS 20*, 2008.
- Fukumizu, K., Bach, F., and Jordan, M. Kernel dimension reduction in regression. *The Annals of Statistics*, 37(4), 2009.
- Fukumizu, K., Song, L., and Gretton, A. Kernel Bayes’ rule. Technical report, arXiv:1009.5736v4, 2011.
- Grünwälder, S., Lever, G., Baldassarre, L., Pontil, M., and Gretton, A. Modelling transition dynamics in mdps with rkhs embeddings. In *ICML*, 2012.
- Kallenberg, O. *Foundations of Modern Probability*. Springer, 2nd edition, 2001.
- Micchelli, C. A. and Pontil, M. On learning vector-valued functions. *Neural Computation*, 17(1):177–204, 2005.
- Shawe-Taylor, J. and Cristianini, N. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- Smola, A., Gretton, A., Song, L., and Schölkopf, B. A Hilbert space embedding for distributions. In *ALT*. Springer, 2007.
- Song, L., Huang, J., Smola, A. J., and Fukumizu, K. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *ICML*, pp. 121, 2009.
- Song, L., Gretton, A., and Guestrin, C. Nonparametric tree graphical models. *AISTATS*, 9, 2010.
- Steinwart, I. and Christmann, A. *Support Vector Machines*. Springer, 2008.
- Werner, D. *Funktionalanalysis*. Springer, 4th edition, 2002.

SUPPLEMENTARY

A. Similarity of minimisers

We assume that all $h \in \mathcal{V}$ are integrable with respect to the regular conditional probability $P(\cdot|x)$ and all $\mu \in \mathcal{H}_\Gamma$ are integrable with respect to \tilde{P} . In particular, these are fulfilled if our general assumptions from Section B.1 are fulfilled.

Lemma A.1. *If there exists $\mu^* \in \mathcal{H}_\Gamma$ such that for any $h \in \mathcal{V}$: $\mathbb{E}[h|X] = \langle h, \mu^*(X) \rangle_{\mathcal{V}}$ $P_{\mathcal{X}}$ -a.s., then for any $\mu \in \mathcal{H}_\Gamma$:*

- (i) $\mathbb{E}_{X,Y} \langle L(Y, \cdot), \mu^*(X) \rangle_{\mathcal{V}} = \mathbb{E}_X \|\mu^*(X)\|_{\mathcal{V}}^2$,
- (ii) $\mathbb{E}_{X,Y} \langle L(Y, \cdot), \mu(X) \rangle_{\mathcal{V}} = \mathbb{E}_X \langle \mu^*(X), \mu(X) \rangle_{\mathcal{V}}$.

Proof. (i) follows from (ii) by setting $\mu := \mu^*$. (ii) can be derived like that:

$$\mathbb{E}_X \langle \mu(X), \mu^*(X) \rangle_{\mathcal{V}} = \mathbb{E}_X \mathbb{E}_Y [\mu(X)(Y)|X] = \mathbb{E}_{X,Y} \langle L(Y, \cdot), \mu(X) \rangle_{\mathcal{V}},$$

where we used Fubini's theorem for the last equality (Kallenberg, 2001, Thm. 1.27). □

Theorem A.2 (Thm 3.1). *If there exists a $\mu^* \in \mathcal{H}_\Gamma$ such that for any $h \in \mathcal{V}$: $\mathbb{E}[h|X] = \langle h, \mu^*(X) \rangle_{\mathcal{V}}$ $P_{\mathcal{X}}$ -a.s., then*

$$\mu^* = \underset{\mu \in \mathcal{H}_\Gamma}{\operatorname{argmin}} \mathcal{E}[\mu] = \underset{\mu \in \mathcal{H}_\Gamma}{\operatorname{argmin}} \mathcal{E}_s[\mu] \quad P_{\mathcal{X}} \text{ a.s.}$$

Furthermore, the minimiser is $P_{\mathcal{X}}$ -a.s. unique.

Proof. We start by showing that the right side is minimised by μ^* . Let μ be any element in \mathcal{H}_Γ then we have

$$\begin{aligned} & \mathbb{E}_{X,Y} [\|L(Y, \cdot) - \mu(X)\|_{\mathcal{V}}^2] - \mathbb{E}_{X,Y} [\|L(Y, \cdot) - \mu^*(X)\|_{\mathcal{V}}^2] \\ &= \mathbb{E}_X \|\mu(X)\|_{\mathcal{V}}^2 - 2\mathbb{E}_{X,Y} \langle L(Y, \cdot), \mu(X) \rangle_{\mathcal{V}} + 2\mathbb{E}_{X,Y} \langle L(Y, \cdot), \mu^*(X) \rangle_{\mathcal{V}} - \mathbb{E}_X \|\mu^*(X)\|_{\mathcal{V}}^2 \\ &= \mathbb{E}_X \|\mu(X)\|_{\mathcal{V}}^2 - 2\mathbb{E}_X \langle \mu^*(X), \mu(X) \rangle_{\mathcal{V}} + \mathbb{E}_X \|\mu^*(X)\|_{\mathcal{V}}^2 = \mathbb{E}_X [\|\mu(X) - \mu^*(X)\|_{\mathcal{V}}^2] \geq 0. \end{aligned}$$

Hence, μ^* is a minimiser. The minimiser is furthermore $P_{\mathcal{X}}$ -a.s. unique: Assume there is a second minimiser μ' then above calculation shows that $0 = \mathbb{E}_{X,Y} [\|L(Y, \cdot) - \mu'(X)\|_{\mathcal{V}}^2] - \mathbb{E}_{X,Y} [\|L(Y, \cdot) - \mu^*(X)\|_{\mathcal{V}}^2] = \mathbb{E}_X [\|\mu^*(X) - \mu'(X)\|_{\mathcal{V}}^2] = \mathbb{E}_X [\|\mu^*(X) - \mu'(X)\|_{\mathcal{V}}^2]$. Hence, $\|\mu^*(X) - \mu'(X)\|_{\mathcal{V}} = 0$ $P_{\mathcal{X}}$ -a.s. (Fremlin, 2000)[122Rc], i.e. a measurable set M with $P_{\mathcal{X}}M = 1$ exists such that $\|\mu^*(X) - \mu'(X)\|_{\mathcal{V}} = 0$ holds for all $X \in M$. As $\|\cdot\|_{\mathcal{V}}$ is a norm we have that $\mu^*(X) = \mu'(X)$ $P_{\mathcal{X}}$ -a.s..

To show the second equality we note that, for every $h \in \mathcal{V}$, $\mathbb{E}_X(\mathbb{E}[h|X] - \langle h, \mu^*(X) \rangle_{\mathcal{V}})^2 = 0$ by assumption. Hence, the supremum over $\mathcal{B}(\mathcal{V})$ is also 0 and the minimum is attained at $\mu^*(X)$. Uniqueness can be seen in the following way: assume there is a second minimiser μ' then for all $h \in \mathcal{H}_{\mathcal{V}}$ we have $\mathbb{E}_X(\langle h, \mu'(X) - \mu^*(X) \rangle_{\mathcal{V}})^2 \leq \mathbb{E}_X(\langle h, \mu'(X) \rangle_{\mathcal{V}} - \mathbb{E}[h|X])^2 + \mathbb{E}_X(\mathbb{E}[h|X] - \langle h, \mu^*(X) \rangle_{\mathcal{V}})^2 = 0$. Hence, $\langle h, \mu'(X) - \mu^*(X) \rangle_{\mathcal{V}} = 0$ $P_{\mathcal{X}}$ -a.s. (Fremlin, 2000)[122Rc], i.e. a measurable set M with $P_{\mathcal{X}}M = 1$ exists such that $\langle h, \mu'(X) - \mu^*(X) \rangle_{\mathcal{V}} = 0$ holds for all $X \in M$. Assume that there exists a $X' \in M$ such that $\mu'(X') \neq \mu^*(X')$ then pick $h := \mu'(X') - \mu^*(X')$ and we have $0 = \langle h, \mu'(X') - \mu^*(X') \rangle_{\mathcal{V}} = \|\mu'(X') - \mu^*(X')\|_{\mathcal{V}} > 0$ as $\|\cdot\|_{\mathcal{V}}$ is a norm. Hence, $\mu'(X') = \mu^*(X')$ a.s.. □

Theorem A.3. *If there exists $\eta > 0$ and $\mu^* \in \mathcal{H}_\Gamma$ such that $\sup_{\mu \in \mathcal{B}(\mathcal{H}_\Gamma)} \mathbb{E}_X [\mathbb{E}[\mu(X)|X] - \langle \mu(X), \mu^*(X) \rangle_{\mathcal{V}}]^2 = \eta < \infty$, that μ' is a minimiser of \mathcal{E}_s and $\hat{\mu}$ is an element of \mathcal{H}_Γ with $\mathcal{E}_s[\hat{\mu}] \leq \mathcal{E}_s[\mu'] + \delta$ then*

- (i) $\mathcal{E}[\mu'] \leq \left(\sqrt{\mathcal{E}[\mu^*]} + \eta^{1/4} \sqrt{8(\|\mu^*\|_{\Gamma} + \|\mu'\|_{\Gamma})} \right)^2$
- (ii) $\mathcal{E}[\hat{\mu}] \leq \left(\sqrt{\mathcal{E}[\mu^*]} + \eta^{1/4} \sqrt{8(\|\mu^*\|_{\Gamma} + \|\hat{\mu}\|_{\Gamma})} + \delta^{1/2} \right)^2$.

Proof. First, observe that if $\mu \in \mathcal{H}_\Gamma$ then we have due to the Jensen inequality

$$\begin{aligned} & |\mathbb{E}_X \mathbb{E}[\mu(X)|X] - \mathbb{E}_X \langle \mu(X), \mu^*(X) \rangle_{\mathcal{V}}| \leq \|\mu\|_{\Gamma} \mathbb{E}_X \left| \mathbb{E} \left[\frac{\mu(X)}{\|\mu\|_{\Gamma}} |X \right] - \left\langle \frac{\mu(X)}{\|\mu\|_{\Gamma}}, \mu^*(X) \right\rangle_{\mathcal{V}} \right| \\ & \leq \|\mu\|_{\Gamma} \sqrt{\mathbb{E}_X \left(\mathbb{E} \left[\frac{\mu(X)}{\|\mu\|_{\Gamma}} |X \right] - \left\langle \frac{\mu(X)}{\|\mu\|_{\Gamma}}, \mu^*(X) \right\rangle_{\mathcal{V}} \right)^2} = \|\mu\|_{\Gamma} \sqrt{\eta}. \end{aligned}$$

We can now reproduce the proof of Lemma A.1 with an approximation error. For any $\mu \in \mathcal{H}_\Gamma$ we have

$$|\mathbb{E}_X \langle \mu(X), \mu^*(X) \rangle_{\mathcal{V}} - \mathbb{E}_{X,Y} \langle L(Y, \cdot), \mu(X) \rangle_{\mathcal{V}}| = |\mathbb{E}_X \langle \mu(X), \mu^*(X) \rangle_{\mathcal{V}} - \mathbb{E}_X \mathbb{E}[\mu(X)|X]| \leq \|\mu\|_\Gamma \sqrt{\eta}.$$

In particular,

$$|\mathbb{E}_{X,Y} \langle L(y, \cdot), \mu^*(X) \rangle_{\mathcal{V}} - \mathbb{E}_{X,Y} \|\mu^*(X)\|_{\mathcal{V}}^2| \leq \|\mu^*\|_\Gamma \sqrt{\eta}.$$

Like in the proof of Thm. 3.1 we have for any μ that

$$\begin{aligned} \mathcal{E}_s[\mu] - \mathcal{E}_s[\mu^*] &= \mathbb{E}_{X,Y} [\|L(Y, \cdot) - \mu(X)\|_{\mathcal{V}}^2] - \mathbb{E}_{X,Y} [\|L(Y, \cdot) - \mu^*(X)\|_{\mathcal{V}}^2] \\ &\geq \mathbb{E}_X \|\mu(X)\|_{\mathcal{V}}^2 - 2\mathbb{E}_X \langle \mu^*(X), \mu(X) \rangle_{\mathcal{V}} + \mathbb{E}_X \|\mu^*(X)\|_{\mathcal{V}}^2 - 2\|\mu^*\|_\Gamma \sqrt{\eta} - 2\|\mu\|_\Gamma \sqrt{\eta} \\ &= \mathbb{E}_X [\|\mu^*(X) - \mu(X)\|_{\mathcal{V}}^2] - 2\sqrt{\eta}(\|\mu^*\|_\Gamma + \|\mu\|_\Gamma). \end{aligned} \quad (18)$$

In particular, $|\mathcal{E}_s[\mu] - \mathcal{E}_s[\mu^*]| \geq \mathcal{E}_s[\mu] - \mathcal{E}_s[\mu^*] \geq \mathbb{E}_X [\|\mu^*(X) - \mu(X)\|_{\mathcal{V}}^2] - 2\sqrt{\eta}(\|\mu^*\|_\Gamma + \|\mu\|_\Gamma)$ and hence

$$\mathbb{E}_X [\|\mu^*(X) - \mu(X)\|_{\mathcal{V}}^2] \leq |\mathcal{E}_s[\mu] - \mathcal{E}_s[\mu^*]| + 2\sqrt{\eta}(\|\mu^*\|_\Gamma + \|\mu\|_\Gamma). \quad (19)$$

We can now bound the error $\mathcal{E}[\mu]$ in dependence of how similar μ is to μ^* in the surrogate cost function \mathcal{E}_s :

$$\begin{aligned} \sqrt{\mathcal{E}[\mu]} &= \sqrt{\mathcal{E}[\mu^*]} + \sup_{h \in \mathcal{B}(\mathcal{V})} \sqrt{\mathbb{E}_X [\langle h, \mu^*(X) - \mu(X) \rangle_{\mathcal{V}}]^2} \leq \sqrt{\mathcal{E}[\mu^*]} + \sqrt{\mathbb{E}_X \left(\frac{\|\mu^*(X) - \mu(X)\|_{\mathcal{V}}^2}{\|\mu^*(X) - \mu(X)\|_{\mathcal{V}}} \right)^2} \\ &\leq \sqrt{\mathcal{E}[\mu^*]} + \eta^{1/4} \sqrt{2(\|\mu^*\|_\Gamma + \|\mu\|_\Gamma)} + \sqrt{|\mathcal{E}_s[\mu] - \mathcal{E}_s[\mu^*]|}, \end{aligned} \quad (20)$$

where we used that $\langle \frac{\mu^*(X) - \mu(X)}{\|\mu^*(X) - \mu(X)\|_{\mathcal{V}}}, \mu^*(X) - \mu(X) \rangle_{\mathcal{V}} \geq \langle h, \mu^*(X) - \mu(X) \rangle_{\mathcal{V}}$ for any $h \in \mathcal{B}(\mathcal{V})$ (observe that h is independent of X) and eq. 19.

Now, for $\mu := \mu'$ observe that $\mathcal{E}_s[\mu'] + 2\sqrt{\eta}(\|\mu^*\|_\Gamma + \|\mu\|_\Gamma) \geq \mathcal{E}_s[\mu^*]$ follows from eq. (18) and as μ' is a \mathcal{E}_s minimiser we have $|\mathcal{E}[\mu^*] - \mathcal{E}[\mu']| \leq 2\sqrt{\eta}(\|\mu^*\|_\Gamma + \|\mu\|_\Gamma)$ and from eq. 20 we get

$$\sqrt{\mathcal{E}[\mu']} \leq \sqrt{\mathcal{E}[\mu^*]} + \eta^{1/4} \sqrt{8(\|\mu^*\|_\Gamma + \|\mu'\|_\Gamma)}.$$

Similarly, for $\mu := \hat{\mu}$ we have

$$\sqrt{\mathcal{E}[\hat{\mu}]} \leq \sqrt{\mathcal{E}[\mu^*]} + \eta^{1/4} \sqrt{8(\|\mu^*\|_\Gamma + \|\hat{\mu}\|_\Gamma)} + \delta^{1/2}.$$

□

B. Assumptions

B.1. Simple assumptions to make the convergence theorems hold.

In this section, we present general assumptions that make the assumptions in (Caponnetto & De Vito, 2007, Theorems 1 and 2) hold. Weaker assumptions are certainly possible and we do not claim that our set of assumptions are the most general possible.

B.1.1. ASSUMPTIONS FOR THE INVOLVED SPACES AND KERNELS

We assume that \mathcal{X} is a compact subset of \mathbb{R}^n , \mathcal{Y} is a compact set and \mathcal{V} is a finite dimensional RKHS. Furthermore, we assume that we work on measure spaces $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$, $(\mathcal{V}, \mathcal{B}(\mathcal{V}))$, where $\mathcal{B}(\mathbb{R})$ etc. are the respective Borel algebras and each of the spaces is equipped with its norm induced topology. Finally, we assume that Γ and L are continuous. Hence, $L(y, y) \leq B < \infty$ for all $y \in \mathcal{Y}$.

B.1.2. DATA GENERATING DISTRIBUTION

The main object in this study is the probability measure \tilde{P} on the space $\mathcal{X} \times \mathcal{Y}$ from which the data is generated. Formally, we have a measure space² $(\mathcal{X} \times \mathcal{Y}, \tilde{\Sigma}, \tilde{P})$ where $\tilde{\Sigma}$ is a suitable σ -algebra and $\tilde{P} : \tilde{\Sigma} \mapsto [0, 1]$ a probability measure. Now we have a transformation from $\phi : y \mapsto L(y, \cdot)$ which turns a sample $\{(x_1, y_1), (x_2, y_2), \dots\}$ into $\{(x_1, L(y_1, \cdot)), (x_2, L(y_2, \cdot)), \dots\}$ with $(x_i, L(y_i, \cdot)) \in \mathcal{X} \times \mathcal{V}$. The measure \tilde{P} and the σ -algebra $\tilde{\Sigma}$ induce through the image measure construction a measure space $(\mathcal{X} \times \mathcal{V}, \Sigma, P)$, where $\Sigma := \{E | \phi^{-1}[E] \in \tilde{\Sigma}, E \subset \mathcal{X} \times \mathcal{V}\}$ is the largest σ -algebra on $\mathcal{X} \times \mathcal{V}$ such that ϕ is measurable and $PE = \tilde{P}(\phi^{-1}[E])$ for all $E \in \Sigma$ (Fremlin, 2000)[112E].

We assumed already that \mathcal{X}, \mathcal{V} carry the Borel algebra. In addition assume that $\tilde{\Sigma} = \mathcal{B}(\mathcal{X} \times \mathcal{Y})$ and $\Sigma = \mathcal{B}(\mathcal{X} \times \mathcal{V})$, i.e. the Borel algebra on the product spaces. The assumptions we stated till now guarantee us that the regular conditional probabilities $P(B|x)$ and $\tilde{P}(\tilde{B}|x)$ exists, where $B \in \mathcal{B}(\mathcal{V})$ and $\tilde{B} \in \mathcal{B}(\mathcal{Y})$ (Steinwart & Christmann, 2008)[Lem. A.3.16].

Our assumptions on L and Γ guarantee that all h in \mathcal{V} and $\mu \in \mathcal{H}_\Gamma$ are continuous (Berlinet & Thomas-Agnan, 2004)[Th. 17] and Cor. 2.3. Hence, all h and μ are measurable as by assumption all spaces are equipped with Borel algebras (Fremlin, 2003)[4A3Cb,d]. Furthermore, as each $|h|$ and $|\mu|$ is a continuous function on a compact set it is upper bounded by a constant B . Thus all h and μ are integrable with respect to any probability measure on $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$ and $(\mathcal{X} \times \mathcal{V}, \mathcal{B}(\mathcal{X} \times \mathcal{V}))$.

The final assumption is that one of the two equivalent conditions of theorem B.3 is fulfilled.

B.2. Verification of the Caponnetto assumptions

Theorem 1-2 from (Caponnetto & De Vito, 2007) are based on a set of assumptions. The assumptions can be grouped into three categories: first, assumptions about the space \mathcal{X} and \mathcal{V} , second assumptions about the space of regression functions that we use and third about the underlying true probability distribution. Table 1 summarizes the assumptions that need to be fulfilled in our setting. In the following we discuss each assumption shortly and we show ways to verify them. Also note, that the general assumption from the last section guarantee that all assumptions in table 1 are fulfilled.

Assumptions for \mathcal{X} and \mathcal{V} \mathcal{X} needs to be a Polish space, that is a separable completely metrizable topological space. If \mathcal{X} is, for example, a subset of \mathbb{R}^n equipped with any metric then \mathcal{X} is Polish if \mathcal{X} is closed.

\mathcal{V} needs to be a separable Hilbert space to be able to apply the theorems. In our case \mathcal{V} is an RKHS so certainly a Hilbert space. Not all RKHSs are separable, but, due to our assumption that \mathcal{V} is finite dimensional it follows that it is separable (e.g. take a basis e_1, \dots, e_n of \mathcal{V} then the countable set $\{\sum_{i \leq n} q_i e_i | q_i \in \mathbb{Q}\}$ is dense in the RKHS norm).

Finally, $\text{Tr}(\Gamma(x, x)) \leq B$ needs to hold for all $x \in \mathcal{X}$, i.e. the trace needs to be bounded over \mathcal{X} .

Assumptions for \mathcal{H}_Γ \mathcal{H}_Γ needs to be separable. Separability of \mathcal{H}_Γ and continuity of the kernel are closely related in the scalar case (Berlinet & Thomas-Agnan, 2004)[Sec. 1.5]. Similarly, we have:

Theorem B.1. *\mathcal{H}_Γ is separable if \mathcal{X} is a Polish space, \mathcal{V} is a finite dimensional Hilbert space and Γ is continuous in each argument with respect to the topology on \mathcal{X} .*

Corollary B.2. *If $\mathcal{X} = \mathbb{R}^n$, \mathcal{V} is a finite dimensional RKHS and Γ is a continuous kernel then \mathcal{H}_Γ is separable.*

Furthermore, the point evaluator $\Gamma_x^*(f) = f(x)$, $f \in \mathcal{H}_\Gamma$, $\Gamma_x^* : \mathcal{H}_\Gamma \mapsto \mathcal{V}$ needs to be Hilbert-Schmidt for every x . This is in our case always fulfilled as \mathcal{V} is finite dimensional: Let $\{e_i\}_{i \leq n}$ be an orthonormal basis of \mathcal{V} :

$$\|\Gamma_x^*\|_{HS}^2 = \|\Gamma_x\|_{HS}^2 = \sum_{i \leq n} \|\Gamma_x e_i\|_\Gamma^2 = \sum_{i \leq n} \langle \Gamma_x e_i, \Gamma_x e_i \rangle_\Gamma = \sum_{i \leq n} \langle e_i, \Gamma(x, x) e_i \rangle_\mathcal{V}.$$

But, $\Gamma(x, x)$ is a finite dimensional matrix as \mathcal{V} is finite dimensional and hence the sum is finite and Γ_x has finite Hilbert-Schmidt norm.

The next condition concerns the measurability of the kernel, $\forall f, g \in \mathcal{V}, B : (x, x') \rightarrow \langle f, \Gamma(x', x) g \rangle_\mathcal{V}$ needs to be measurable. The simplest case where this holds is where Γ is continuous as a map from (y, x) to $\mathcal{L}(\mathcal{V})$ and where we equip \mathbb{R}, \mathcal{X} and \mathcal{V} with the Borel algebra. Then we have measurability as compositions from continuous functions are continuous and continuous functions are Borel measurable.

²Some measure theory is needed to write things down cleanly. See also (Steinwart & Christmann, 2008)[chp. 2] where similar problems are discussed.

B.3. Assumptions for the “true” probability measure

There are further assumption concerning the measures P and \tilde{P} from which the data is generated that need to be fulfilled to be able to apply the convergence theorems. We now follow up on the discussion of these measures in appendix B.1.2. Again our general assumptions are sufficient to guarantee that the assumptions are fulfilled.

The first property that P needs to fulfill is that $\int_{\mathcal{X} \times \mathcal{Y}} \|h\|_{\mathcal{V}}^2 dP(x, h) < \infty$ where the integral is the Lebesgue integral. This is essentially trivially fulfilled for all P in our setting as P is concentrated on elements $L(y, \cdot)$, i.e. $\mathcal{L} = \{L(y, \cdot) | y \in \mathcal{Y}\} \in \Sigma$ as $\phi^{-1}[\mathcal{L}] = \mathcal{X} \times \mathcal{Y} \in \tilde{\Sigma}$ and $P\mathcal{L} = \tilde{P}(\phi^{-1}\mathcal{L}) = \tilde{P}(\mathcal{X} \times \mathcal{Y}) = 1$. Furthermore, $\|L(y, \cdot)\|_{\mathcal{V}} = L(y, y) < B < \infty$ if we have a bounded kernel L . The only problem is that we need to guarantee that $\|h\|_{\mathcal{V}}^2$ is integrable. Let us therefore assume in the following that $\|h\|_{\mathcal{V}}^2 \upharpoonright \mathcal{L}$ (i.e. the restriction to \mathcal{L}) is measurable. We have $\|h\|_{\mathcal{V}}^2 \upharpoonright \mathcal{L} < B\chi_{\mathcal{L}}$ where χ is the characteristic function. $B\chi_{\mathcal{L}}$ is measurable as \mathcal{L} is measurable. Furthermore, it is integrable as it is a simple function and $P\mathcal{L} = 1 < \infty$. Now, $\|h \upharpoonright \mathcal{L}\|_{\mathcal{V}}^2$ is defined on the conegible set \mathcal{L} (i.e. $P\mathcal{L} = 1$) and by assumption $\|h\|_{\mathcal{V}}^2 \upharpoonright \mathcal{L}$ is measurable. Hence, $\|h\|_{\mathcal{V}}^2 \upharpoonright \mathcal{L}$ is integrable (Fremlin, 2000)[Lem. 122Jb] and as $\|h\|_{\mathcal{V}}^2 \upharpoonright \mathcal{L} =_{P-a.e.} \|h\|_{\mathcal{V}}^2$ we have that $\|h\|_{\mathcal{V}}^2$ is integrable (Fremlin, 2000)[122Ld] with integral $\int_{\mathcal{X} \times \mathcal{Y}} \|h\|_{\mathcal{V}}^2 dP(x, v) \leq B < \infty$.

The next assumption concerns the infimum $\mathcal{E}[\mu^*] = \inf_{\mu \in \mathcal{H}_{\Gamma}} \mathcal{E}[\mu]$, i.e. the infimum must be attained at a $\mu^* \in \mathcal{H}_{\Gamma}$. One can restate this condition in terms of the RKHS norm of a sequence converging to the infimum: the infimum is attained iff such a sequence exists which is bounded in the RKHS norm. The intuition is here that the condition is not fulfilled if we need to make μ_n more and more complex (in the sense of a high RKHS norm) to optimize the risk.

Theorem B.3. *Let $\|h\|_{\mathcal{V}}^2, \|\mu(x) - h\|_{\mathcal{V}}$ be integrable for all $\mu \in \mathcal{H}_{\Gamma}$ and let \mathcal{V} be a finite dimensional Hilbert space. Then there exists a $\mu^* \in \mathcal{H}_{\Gamma}$ with $\mathcal{E}[\mu^*] = \inf_{\mu \in \mathcal{H}_{\Gamma}} \mathcal{E}[\mu]$ iff a $B > 0$ exists and a sequence $\{\mu_n\}_{n \in \mathbb{N}}$ with $\mathcal{E}[\mu_n] \leq \inf_{\mu \in \mathcal{H}_{\Gamma}} \mathcal{E}[\mu] + 1/n$ and $\|\mu_n\|_{\Gamma} < B$.*

Remark The important assumption is that the sequence $\{\mu_n\}$ is bounded in the RKHS norm as we can by definition always find a sequence $\{\mu_n\}$ that converges to the infimum.

Proof. “ \Rightarrow ”: set $\mu_n := \mu^*$ then obviously the sequence converges to the infimum and is bounded as $\|\mu_n\|_{\Gamma} = \|\mu^*\|_{\Gamma} < \infty$ as $\mu^* \in \mathcal{H}_{\Gamma}$.

“ \Leftarrow ”: We need to pull the limit through the integral and we need to make sure that the limit is attained by an RKHS function. First, let us check for the limit to be in the RKHS. $\{\mu_n\}_{n \in \mathbb{N}}$ is a bounded sequence in a reflexive space (every Hilbert space is reflexive). Hence, there is a sub sequence $\{\mu_{n_k}\}_{k \in \mathbb{N}}$ which converges weakly (Werner, 2002)[Th. III.3.7]. Furthermore, \mathcal{H}_{Γ} is closed and convex. Therefore, the weak limit μ is in \mathcal{H}_{Γ} (Werner, 2002)[Thm III.3.8].

Now we turn to the integral. We want to apply the Lebesgue dominated convergence theorem to move the limit inside the integral. We need to construct a suitable sequence: Consider the sequence $s_k(x, h) := \{\|g_k(x) - h\|_{\mathcal{V}}\}_{k \in \mathbb{N}}$ in \mathbb{R} , where $g_k := \mu_{n_k}$ and $\{\mu_{n_k}\}_{k \in \mathbb{N}}$ is the convergent sub sequence. The sequence $\{s_k(x, v)\}_{k \in \mathbb{N}}$ is a Cauchy sequence in \mathbb{R} as for any $\epsilon > 0$ we have

$$\left| \|g_k(x) - h\|_{\mathcal{V}} - \|g_l(x) - h\|_{\mathcal{V}} \right| \leq \|g_k(x) - g_l(x)\|_{\mathcal{V}} = \sum_{i \leq u} \langle g_k(x) - g_l(x), e_i \rangle_{\mathcal{V}} = \sum_{i \leq u} \langle g_k - g_l, \Gamma_x e_i \rangle_{\Gamma},$$

where we used the Parseval equality and set u as the dimension of \mathcal{V} . Now, as $g_k(x)$ is weakly convergent we can find for each i a $N(i)$ such that $\langle g_k - g_l, \Gamma_x e_i \rangle_{\Gamma} < \epsilon/u$. Let $N = \max\{N(1), \dots, N(u)\}$ then we have

$$\left| \|g_k(x) - h\|_{\mathcal{V}} - \|g_l(x) - h\|_{\mathcal{V}} \right| < \epsilon.$$

As \mathbb{R} is complete the sequence $\{s_k(x, h)\}_{k \in \mathbb{N}}$ attains its limit for every pair $(x, h) \in \bigcup_{l \in \mathbb{N}} \bigcap_{k \geq l} \text{dom}(\|g_k(x) - h\|_{\mathcal{V}})$, i.e. whenever the sequence is defined for a pair (x, h) . In particular, as $\|g_k(x) - h\|_{\mathcal{V}}$ is integrable there exists for every k a measurable set $A_k \subset \text{dom}(\|g_k(x) - h\|_{\mathcal{V}})$ with $P[A_k] = 1$. s_k attains its limit on the set $D := \bigcup_{l \in \mathbb{N}} \bigcap_{k \geq l} A_k \in \Sigma$ and $P[D] = 1$ (continuity of measure). Finally, $\|g_k(x) - h\|_{\mathcal{V}} \leq B\chi_D + \|h\|_{\mathcal{V}} \leq (B+1)\chi_D + \|h\|_{\mathcal{V}}^2 =: f(x, h)$. $f(x, h)$ is integrable as a sum of integrable functions. Hence, we can apply the Lebesgue dominated convergence theorem (Fremlin, 2000)[Th. 123C] and we get that $\lim_{k \rightarrow \infty} \|g_k(x) - h\|_{\mathcal{V}}$ is integrable and

$$\lim_{k \rightarrow \infty} \mathcal{E}[g_k] = \int_D \lim_{k \rightarrow \infty} \|g_k(x) - h\|_{\mathcal{V}} dP(x, h).$$

Furthermore, $\lim_{k \rightarrow \infty} \|g_k(x) - h\|_{\mathcal{V}} = \|\mu(x) - h\|_{\mathcal{V}}$ on D as for any $\epsilon > 0$ there exists a N such that for all $k > N$

$$\left| \|g_k(x) - h\|_{\mathcal{V}} - \|\mu(x) - h\|_{\mathcal{V}} \right| \leq \|g_k(x) - \mu(x)\|_{\mathcal{V}} \leq \sum_{i \leq u} \langle g_k(x) - \mu(x), e_i \rangle_{\mathcal{V}} = \sum_{i \leq u} \langle g_k - \mu, \Gamma_x e_i \rangle_{\Gamma} < \epsilon$$

as g_k converges weakly to μ and $u < \infty$. Hence, $\lim_{k \rightarrow \infty} \mathcal{E}[g_k] = \mathcal{E}[\mu]$ and $\mu \in \mathcal{H}_{\Gamma}$. Also,

$$\inf_{g \in \mathcal{H}_{\Gamma}} \mathcal{E}[g] \leq \mathcal{E}[\mu] \leq \inf_{g \in \mathcal{H}_{\Gamma}} \mathcal{E}[g] + 1/n$$

for all $n \in \mathbb{N}$. Therefore, $\inf_{g \in \mathcal{H}_{\Gamma}} \mathcal{E}[g] = \mathcal{E}[\mu]$. Hence, $\mu^* := \mu$ is a minimiser. □

Furthermore, the probability measure needs to factor as $P(h, x) = P_X(x)P(h|x)$ and there needs to be constants c, d such that

$$\int_{\mathcal{V}} \left(e^{\frac{\|h - \mu^*(x)\|_{\mathcal{V}}}{d}} - \frac{\|h - \mu^*(x)\|_{\mathcal{V}}}{d} - 1 \right) dP(h|x) \leq \frac{c^2}{2d^2}.$$

In particular, $\exp\left(\frac{\|h - \mu^*(x)\|_{\mathcal{V}}}{d}\right) - \frac{\|h - \mu^*(x)\|_{\mathcal{V}}}{d} - 1$ needs to be $P(\cdot|x)$ integrable. In case that $P(h|x)$ is concentrated on $\mathcal{L} := \{L(x, \cdot) | x \in \mathcal{X}\}$ we have $\exp\left(\frac{\|h - \mu^*(x)\|_{\mathcal{V}}}{d}\right) - \frac{\|h - \mu^*\|_{\mathcal{V}}}{d} - 1 \leq \exp\left(\frac{L(y, y) + B\|\mu^*\|_{\Gamma}}{d}\right)$ on \mathcal{L} , where we used Corollary 4 and restricted $d > 0$. As μ^* has finite RKHS norm and as we have a probability measure the assumption is always fulfilled in our case.

B.4. Discussion of assumptions from (Song et al., 2009)

We review conditions from (Song et al., 2009; Fukumizu et al., 2011) that establish the existence of $\mu \in \mathcal{H}_{\Gamma}$ such that $\mathbb{E}(h(Y)|X = x) = \langle h, \mu(x) \rangle_{\mathcal{H}_L} \forall h \in \mathcal{H}_L$, as well as convergence results for an empirical estimate of μ (note that certain additional conditions over those stated in (Fukumizu et al., 2011) are needed for the results to be rigorous, following the discussion in (Fukumizu et al., 2011, Appendix, Theorem 6.6): we include these here). To describe these conditions, some preliminary definitions will be needed. Define the uncentred covariance operators,

$$C_{XX} = \mathbb{E}_X K(X, \cdot) \otimes K(X, \cdot) \quad C_{XY} = \mathbb{E}_{XY} K(X, \cdot) \otimes L(X, \cdot),$$

where the tensor product satisfies $(a \otimes b)c = \langle b, c \rangle_{\mathcal{H}_K} a$ for $a \in \mathcal{H}_K$ and $b, c \in \mathcal{H}_L$ (an analogous result applies for $a, b, c \in \mathcal{H}_K$). As the mappings in Song et al. are defined from \mathcal{H}_K to \mathcal{H}_L , rather than \mathcal{X} to \mathcal{H}_L , we next establish the relation between the Song et al. notation and ours: for every $h \in \mathcal{H}_L$,

$$\langle h, \mu(x) \rangle_{\mathcal{H}_L} = \langle h, \mathcal{U}_{Y|X} \phi(x) \rangle_{\mathcal{H}_L},$$

where the notation on the left is from the present paper, that on the right is from Song et al., and $\mathcal{U}_{Y|X} : \mathcal{H}_K \rightarrow \mathcal{H}_L$ (we therefore emphasize $\mathcal{U}_{Y|X} \notin \mathcal{H}_{\Gamma}$, since its input space is \mathcal{H}_K , however $\mathcal{U}_{Y|X} \circ \phi(x) \in \mathcal{H}_{\Gamma}$ under the conditions stated below). We have the following theorem, corresponding to (Song et al., 2009, Theorem 1).

Theorem B.4 ((Song et al., 2009)). *Assume that for all $h \in \mathcal{H}_L$, we have $\mathbb{E}(h(Y)|X = \cdot) \in \mathcal{H}_K$, C_{XX} is injective, and the operator $Sh := \mathbb{E}(h(Y)|X = \cdot)$ is bounded for all $h \in \mathcal{H}_L$. Song et al. make the further smoothness assumption that $C_{XX}^{-3/2} C_{XY}$ is Hilbert-Schmidt (in fact, following (Fukumizu et al., 2011, Appendix), boundedness should be sufficient, although this is outside the scope of the present paper). Then the definition $\mathcal{U}_{Y|X} := C_{YX} C_{XX}^{-1}$ satisfies both $\mathcal{U}_{Y|X} \circ \phi(x) \in \mathcal{H}_{\Gamma}$ and $\mu(x) = \mathcal{U}_{Y|X} \phi(x)$, and an empirical estimate $\hat{\mathcal{U}}_{Y|X} := \hat{C}_{YX} \left(\hat{C}_{XX} + \delta_n I \right)^{-1}$ converges with rate*

$$\left\| \hat{C}_{YX} \left(\hat{C}_{XX} + \delta_n I \right)^{-1} \phi(x) - C_{YX} C_{XX}^{-1} \phi(x) \right\|_{\mathcal{G}}^2 = O_p(n^{-1/4}).$$

when $\delta_n = n^{-1/4}$.

C. Proofs

Convergence result from (Song et al., 2009) The theorem in (Song et al., 2009) states the following: Let P be a probability measure and let \hat{U} be the estimate of U defined in (Song et al., 2009) then there exists a constant C such that $\lim_{n \rightarrow \infty} P^n[\|U - \hat{U}\|_{HS} > Cn^{-1/8}] = 0$, where HS denotes Hilbert-Schmidt norm.

Under the assumption that $K(x, x) < B$ for all $x \in \mathcal{X}$ this implies: for a probability measure P and an estimate \hat{U} of U it holds that there exists a constant τ such that $\lim_{n \rightarrow \infty} P^n[\sup_{h \in \mathcal{B}(\mathcal{V})} \mathbb{E}_X[\mathbb{E}[h|X] - \langle h, \hat{\mu}(X) \rangle_{\mathcal{V}}]^2 > \tau n^{-1/4}] = 0$.

Proof. Let $\|\cdot\|$ be the operator norm. We have that $\|A\| \leq \|A\|_{HS}$ for a Hilbert-Schmidt operator (Werner, 2002)[p.268,Satz VI.6.2c]. Using this fact and Cauchy-Schwarz we get:

$$\begin{aligned} \sup_{h \in \mathcal{B}(\mathcal{V})} \mathbb{E}_X[\mathbb{E}[h|X] - \langle h, \hat{U}K(X, \cdot) \rangle_{\mathcal{V}}]^2 &\leq \mathbb{E}_X[\|U - \hat{U}\| \frac{K(X, \cdot)}{\|K(X, \cdot)\|_K}]_{\mathcal{V}}^2 \|K(X, \cdot)\|_K^2 \leq \|U - \hat{U}\|^2 \mathbb{E}_X K^2(X, X) \\ &\leq B \|U - \hat{U}\|_{HS}^2 \end{aligned}$$

and the statement follows with $\tau := C^2/B$. \square

Example The space \mathcal{X} together with the reproducing kernel of \mathcal{H}_K can lead to problems for the assumption from (Song et al., 2009):

Corollary C.1. *Let \mathcal{V} be finite dimensional such that a function $\tilde{h} \in \mathcal{V}$ exists with $\tilde{h}(y) \geq \epsilon > 0$ for all $y \in \mathcal{Y}$. Furthermore, let $\mathcal{X} := [-1, 1]$ and the reproducing kernel for \mathcal{H}_K be $K(x, y) = xy$. Then there exists no measure which fulfills the assumption from eq. 13.*

Proof. Assume there is a measure and a measure space $(\mathcal{Y}, \Sigma, \mu)$. We can assume that all $h \in \mathcal{V}$ are integrable as otherwise the assumption fails. We have $\mathbb{E}[\tilde{h}|X = x] > \mathbb{E}[\epsilon \chi_{\mathcal{Y}}|X = x] = \epsilon > 0$. However, by assumption $\mathbb{E}[\tilde{h}|X = x] = \sum_{i=1}^m \langle \tilde{h}, h_i \rangle_L \mathbb{E}[\tilde{h}_i|X = x] = \sum_{i \leq m} \langle \tilde{h}, h_i \rangle_{\mathcal{V}} f_i(x)$ for some $f_i \in \mathcal{H}_K$, where m is the dimension of \mathcal{V} and $\{h_i\}_{i \leq m}$ is an orthonormal basis. $\mathcal{H}_K = \{z \cdot | z \in \mathbb{R}\}$ and hence $\mathbb{E}[\tilde{h}(Y)|X = x] = \langle \tilde{h}, \sum_{i \leq m} z_i x h_i \rangle_{\mathcal{V}} = x \langle \tilde{h}, \sum_{i \leq m} z_i h_i \rangle_{\mathcal{V}}$. If the inner product is positive then we get $\mathbb{E}[\tilde{h}(Y)|X = x] \leq 0$ for $x \in [-1, 0]$ and a contradiction. Similarly for the case where the inner product is negative. \square

D. Justification of sparsity conditional mean embedding

In this section, we provide a justification for the objective function (16) which we used in Section 5 to derive a sparse conditional mean embedding. Specifically, we show that the objective function (16) provides a natural upper bound to (5), which is the error measure we used to derive the embedding itself.

For simplicity, we assume that the underlying conditional mean embedding belongs to a reproducing kernel Hilbert space of \mathcal{H}_L -valued functions, \mathcal{H}_{Γ} , with operator-value kernel $\Gamma \in \mathcal{L}(\mathcal{H}_L)$. That is, there exists $\mu^* \in \mathcal{H}_{\Gamma}$, such that, for every $h \in \mathcal{H}_L$ and $x \in \mathcal{X}$, it holds that $\mathbb{E}_Y[h(Y)|X = x] = \langle \mu^*(x), h \rangle_L$. Under this assumption we have that

$$\begin{aligned} \mathbb{E}_Y[h(Y)|X] - \langle h, \mu(X) \rangle_L &= \langle \mu^*(x) - \mu(x), h \rangle_L \\ &= \langle \mu^* - \mu, \Gamma(x, \cdot) \rangle_{\Gamma} \\ &\leq \|\mu^* - \mu\|_{\Gamma} \|\Gamma(x, \cdot)\|_L \\ &\leq \|\mu^* - \mu\|_{\Gamma} \|\Gamma(x, x)\|_{\Gamma}^{\frac{1}{2}} \|h\|_L. \end{aligned}$$

where the second equality follows from the reproducing property of the kernel Γ , the first inequality follows from Cauchy Schwarz's inequality, the last inequality follows from (Micchelli & Pontil, 2005, Proposition 1(d)) and $\|\Gamma(x, x)\|_{\Gamma}$ denotes the operator norm. Consequently,

$$\begin{aligned} \mathcal{E}[\mu] &= \sup_{\|h\|_L \leq 1} \mathbb{E}_X \left[(\mathbb{E}_Y[h(Y)|X] - \langle h, \mu(X) \rangle_L)^2 \right] \\ &\leq c \|\mu^* - \mu\|_{\Gamma}^2 \end{aligned} \tag{21}$$

where we have defined $c = \mathbb{E}_X[\|\Gamma(X, X)\|_{\Gamma}]$.

We choose μ to be the solution $\hat{\mu}$ of problem (9). As noted at the beginning of Section 5, in many practical circumstances, we wish to approximate $\hat{\mu}$ by a sparse version μ_{sparse} which involves a smaller number of parameters. For this purpose, it is natural to use the error measure $\|\cdot\|_{\Gamma}$. Indeed, equation (21) and the triangle inequality yield that

$$\mathcal{E}[\mu_{\text{sparse}}] \leq c(\|\mu^* - \hat{\mu}\|_{\Gamma} + \|\hat{\mu} - \mu_{\text{sparse}}\|_{\Gamma})^2.$$

Thus if $\hat{\mu}$ estimates well the true embedding μ^* , so will the sparse mean embedding μ_{sparse} .

In the specific case considered in Section 5, $\Gamma(x, x') = K(x, x')Id$, so that $\|\Gamma(x, x)\| = K(x, x)$ and $\|\cdot\|_{\Gamma} = \|\cdot\|_{K \otimes L}$. Furthermore, we choose $\hat{\mu} = \sum_{i,j=1}^n W_{ij} K_{x_i} L_{y_j}$ and $\mu_{\text{sparse}} = \sum_{i,j=1}^n M_{ij} K_{x_i} L_{y_j}$, where matrix M is encouraged to be sparse. Different sparsity methods are discussed in Section 5, such as the the Lasso method (cf. equation (15)).