

Infinite Task Learning in RKHSs

Romain Brault¹, Alex Lambert², Zoltán Szabó³, Maxime Sangnier⁴, et Florence d’Alché-Buc⁵

¹Thales

²Télécom ParisTech

³UPMC

⁴École Polytechnique

⁵Télécom ParisTech

June 3, 2019

Abstract

Machine learning has witnessed tremendous success in solving tasks depending on a single hyperparameter. When considering simultaneously a finite number of tasks, multi-task learning enables one to account for the similarities of the tasks via appropriate regularizers. A step further consists of learning a continuum of tasks for various loss functions. A promising approach, called *Parametric Task Learning*, has paved the way in the continuum setting for affine models and piecewise-linear loss functions. In this work, we introduce a novel approach called *Infinite Task Learning*: its goal is to learn a function whose output is a function over the hyperparameter space. We leverage tools from operator-valued kernels and the associated Vector-Valued Reproducing Kernel Hilbert Space that provide an explicit control over the role of the hyperparameters, and also allows us to consider new type of constraints. We provide generalization guarantees to the suggested scheme and illustrate its efficiency in cost-sensitive classification, quantile regression and density level set estimation.

Mots-clef: méthodes à noyaux, multi-task learning, multi-output prediction.

1 Introduction

Several fundamental problems in machine learning and statistics can be phrased as the minimization of a loss function described by a hyperparameter. The hyperparameter might capture numerous aspects of the problem: (i) the tolerance w.r.t. outliers as the ϵ -insensitivity in Support Vector Regression [VGS97], (ii) importance of smoothness or sparsity such as the weight of the

l_2 -norm in Tikhonov regularization [TA77], l_1 -norm in LASSO [Tib96], or more general structured-sparsity inducing norms [Bac+12], (iii) Density Level-Set Estimation (DLSE), see for example one-class support vector machines One-Class Support Vector Machine (OCSVM, [Sch+00]), (iv) confidence as exemplified by Quantile Regression (QR, [KB78]), or (v) importance of different decisions as implemented by Cost-Sensitive Classification (CSC, [ZE01]).

In various cases including QR, CSC or DLSE, one is interested in solving the parameterized task for several hyperparameter values. Multi-Task Learning [EP04] provides a principled way of benefiting from the relationship between similar tasks while preserving local properties of the algorithms: ν -property in DLSE [GLM13] or quantile property in QR [Tak+06].

A natural extension from the traditional multi-task setting is to provide a prediction tool being able to deal with *any* value of the hyperparameter. In their seminal work, [Tak+13] extended multi-task learning by considering an infinite number of parametrized tasks in a framework called Parametric Task Learning (PTL). Assuming that the loss is piecewise affine in the hyperparameter, the authors are able to get the whole solution path through parametric programming, relying on techniques developed by Hastie et al. [Has+04].

In this paper, we relax the affine model assumption on the tasks as well as the piecewise-linear assumption on the loss, and take a different angle. We propose Infinite Task Learning (ITL) within the framework of function-valued function learning to handle a continuum number of parameterized tasks. For that purpose we leverage tools from operator-valued kernels and the associated Vector-Valued Reproducing Kernel Hilbert Space (vv-

RKHS, [Ped57]). The idea is that the output is a function on the hyperparameters —modelled as scalar-valued Reproducing Kernel Hilbert Space (RKHS)—, which provides an explicit control over the role of the hyperparameters, and also enables us to consider new type of constraints. In the studied framework each task is described by a (scalar-valued) RKHS over the input space which is capable of dealing with nonlinearities. The resulting ITL formulation relying on vv-RKHS specifically encompasses existing multi-task approaches including joint quantile regression [SFd16] or multi-task variants of density level set estimation [GLM13] by encoding a continuum of tasks.

Our **contributions** can be summarized as follows:

- We propose ITL, a novel vv-RKHS-based scheme to learn a continuum of tasks parametrized by a hyperparameter and design new regularizers.
- We prove excess risk bounds on ITL and illustrate its efficiency in quantile regression, cost-sensitive classification, and density level set estimation.

The paper is structured as follows. The ITL problem is defined in Section 2. In Section 3 we detail how the resulting learning problem can be tackled in vv-RKHSs. Excess risk bounds is the focus of Section 4. Numerical results are presented in Section 5. Conclusions are drawn in Section 6.

2 From parameterized to infinite task learning

After introducing a few notations, we gradually define our goal by moving from single parameterized tasks (Section 2.1) to ITL (Section 2.3) through multi-task learning (Section 2.2).

Notations: $\mathbb{1}_S$ is the indicator function of set S . We use the $\sum_{i,j=1}^{n,m}$ shorthand for $\sum_{i=1}^n \sum_{j=1}^m$. $|x|_+ = \max(x, 0)$ denotes positive part. $\mathcal{F}(\mathcal{X}; \mathcal{Y})$ stands for the set of $\mathcal{X} \rightarrow \mathcal{Y}$ functions. Let \mathcal{Z} be Hilbert space and $\mathcal{L}(\mathcal{Z})$ be the space of $\mathcal{Z} \rightarrow \mathcal{Z}$ bounded linear operators. Let $\mathbf{K} : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{L}(\mathcal{Z})$ be an operator-valued kernel, i.e. $\sum_{i,j=1}^n \langle z_i, \mathbf{K}(x_i, x_j) z_j \rangle_{\mathcal{Z}} \geq 0$ for all $n \in \mathbb{N}^*$ and $x_1, \dots, x_n \in \mathcal{X}$ and $z_1, \dots, z_n \in \mathcal{Z}$ and $\mathbf{K}(x, z) = \mathbf{K}(z, x)^*$ for all $x, z \in \mathcal{X}$. \mathbf{K} gives rise to the Vector-Valued Reproducing Kernel Hilbert Space $\mathcal{H}_{\mathbf{K}} = \overline{\text{span}} \{ \mathbf{K}(\cdot, x) z \mid x \in \mathcal{X}, z \in \mathcal{Z} \} \subset \mathcal{F}(\mathcal{X}; \mathcal{Z})$, where $\overline{\text{span}} \{ \cdot \}$ denotes the closure of the linear span of its argument. For further details on vv-RKHS the reader is referred to [Car+10].

2.1 Learning Parameterized Tasks

A *supervised parametrized task* is defined as follows. Let $(X, Y) \in \mathcal{X} \times \mathcal{Y}$ be a random variable with joint distribution $\mathbf{P}_{X, Y}$ which is assumed to be fixed but unknown; we also assume that $\mathcal{Y} \subset \mathbb{R}$. We have access to n independent identically distributed observations called training samples: $\mathcal{S} := ((x_i, y_i))_{i=1}^n \sim \mathbf{P}_{X, Y}^{\otimes n}$. Let Θ be the domain of hyperparameters, and $v_{\theta} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a loss function associated to $\theta \in \Theta$. Let $\mathcal{H} \subset \mathcal{F}(\mathcal{X}; \mathcal{Y})$ denote our hypothesis class; throughout the paper \mathcal{H} is assumed to be a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$. For a given θ , the goal is to estimate the minimizer of the expected risk

$$\mathbf{R}^{\theta}(\mathbf{h}) := \mathbf{E}_{X, Y}[v_{\theta}(Y, \mathbf{h}(X))] \quad (1)$$

over \mathcal{H} , using the training sample \mathcal{S} . This task can be addressed by solving the regularized empirical risk minimization problem

$$\min_{\mathbf{h} \in \mathcal{H}} \mathbf{R}_{\mathcal{S}}^{\theta}(\mathbf{h}) + \Omega(\mathbf{h}), \quad (2)$$

where $\mathbf{R}_{\mathcal{S}}^{\theta}(\mathbf{h}) := \frac{1}{n} \sum_{i=1}^n v_{\theta}(y_i, \mathbf{h}(x_i))$ is the empirical risk and $\Omega : \mathcal{H} \rightarrow \mathbb{R}$ is a regularizer. Below we give three examples.

Quantile Regression: In this setting $\theta \in (0, 1)$. For a given hyperparameter θ , in Quantile Regression the goal is to predict the θ -quantile of the real-valued output conditional distribution $\mathbf{P}_{Y|X}$. The task can be tackled using the pinball loss [KB78] defined in Eq. (3).

$$v_{\theta}(\mathbf{y}, \mathbf{h}(x)) = |\theta - \mathbb{1}_{\mathbb{R}_-}(\mathbf{y} - \mathbf{h}(x))| |\mathbf{y} - \mathbf{h}(x)|, \quad (3)$$

$$\Omega(\mathbf{h}) = \frac{\lambda}{2} \|\mathbf{h}\|_{\mathcal{H}}^2, \quad \lambda > 0.$$

Cost-Sensitive Classification: Our next example considers binary classification ($\mathcal{Y} = \{-1, 1\}$) where a (possibly) different cost is associated with each class, as it is often the case in medical diagnosis. The sign of $\mathbf{h} \in \mathcal{H}$ yields the estimated class and in cost-sensitive classification one takes

$$v_{\theta}(\mathbf{y}, \mathbf{h}(x)) = \left| \frac{1}{2}(\theta + 1) - \mathbb{1}_{\{-1\}}(\mathbf{y}) \right| |1 - \mathbf{y}\mathbf{h}(x)|_+, \quad (4)$$

$$\Omega(\mathbf{h}) = \frac{\lambda}{2} \|\mathbf{h}\|_{\mathcal{H}}^2, \quad \lambda > 0.$$

The $\theta \in [-1, 1]$ hyperparameter captures the trade-off between the importance of correctly classifying the samples having -1 and $+1$ labels. When θ is close to -1 , the obtained \mathbf{h} focuses on classifying well class -1 , and vice-versa. Typically, it is desirable for a physician to choose *a posteriori* the value of the hyperparameter at which he wants to predict. Since this cost can rarely be considered to be fixed, this motivates to learn one model giving access to all hyperparameter values.

Density Level-Set Estimation: Examples of parameterized tasks can also be found in the unsupervised setting. For instance in outlier detection, the goal is to separate outliers from inliers. A classical technique to tackle this task is OCSVM [Sch+00]. OCSVM has a free parameter $\theta \in (0, 1]$, which can be proven to be an upper bound on the fraction of outliers. When using a Gaussian kernel with a bandwidth tending towards zero, OCSVM consistently estimates density level sets [VV06]. This unsupervised learning problem can be empirically described by the minimization of a regularized empirical risk $R_S^\theta(\mathbf{h}, \mathbf{t}) + \Omega(\mathbf{h})$, solved *jointly* over $\mathbf{h} \in \mathcal{H}$ and $\mathbf{t} \in \mathbb{R}$ with

$$v_\theta(\mathbf{t}, \mathbf{h}(x)) = -\mathbf{t} + \frac{1}{\theta} |\mathbf{t} - \mathbf{h}(x)|_+, \quad \Omega(\mathbf{h}) = \frac{1}{2} \|\mathbf{h}\|_{\mathcal{H}}^2.$$

2.2 Solving a Finite Number of Tasks as Multi-Task Learning

In all the aforementioned problems, one is rarely interested in the choice of a single hyperparameter value (θ) and associated risk (R_S^θ), but rather in the joint solution of multiple tasks. The naive approach of solving the different tasks independently can easily lead to inconsistencies. A principled way of solving many parameterized tasks has been cast as a MTL problem [EMP05] which takes into account the similarities between tasks and helps providing consistent solutions. For example it is possible to encode the similarities of the different tasks in MTL through an explicit constraint function [Cil+17]. In the current work, the similarity between tasks is designed in an implicit way through the use of a kernel on the hyperparameters. Moreover, in contrast to MTL, in our case the input space and the training samples are the same for each task; a task is specified by a value of the hyperparameter. This setting is sometimes referred to as multi-output learning [ÁRL12].

Formally, assume that we have p tasks described by parameters $(\theta_j)_{j=1}^p$. The idea of multi-task learning is to minimize the sum of the local loss functions $R_S^{\theta_j}$, i. e.

$$\arg \min_{\mathbf{h}} \sum_{j=1}^p R_S^{\theta_j}(\mathbf{h}_j) + \Omega(\mathbf{h}),$$

where the individual tasks are modelled by the real-valued \mathbf{h}_j functions, the overall \mathbb{R}^p -valued model is the vector-valued function $\mathbf{h}: x \mapsto (\mathbf{h}_1(x), \dots, \mathbf{h}_p(x))$, and Ω is a regularization term encoding similarities between tasks.

It is instructive to consider two concrete examples:

- In joint quantile regression one can use the regularizer to encourage that the predicted conditional quantile

estimates for two similar quantile values are similar. This idea forms the basis of the approach proposed by Sangnier et al. [SFd16] who formulates the joint quantile regression problem in a vector-valued Reproducing Kernel Hilbert Space with an appropriate decomposable kernel that encodes the links between the tasks. The obtained solution shows less quantile curve crossings compared to estimators not exploiting the dependencies of the tasks as well as an improved accuracy.

- A multi-task version of DLSE has recently been presented by Glazer et al. [GLM13] with the goal of obtaining nested density level sets as θ grows. Similarly to joint quantile regression, it is crucial to take into account the similarities of the tasks in the joint model to efficiently solve this problem.

2.3 Towards Infinite Task Learning

In the following, we propose a novel framework called Infinite Task Learning in which we learn a function-valued function $\mathbf{h} \in \mathcal{F}(\mathcal{X}; \mathcal{F}(\Theta; \mathcal{Y}))$. Our goal is to be able to handle new tasks after the learning phase and thus, not to be limited to given predefined values of the hyperparameter. Regarding this goal, our framework generalizes the Parametric Task Learning approach introduced by Takeuchi et al. [Tak+13], by allowing a wider class of models and relaxing the hypothesis of piece-wise linearity of the loss function. Moreover a nice byproduct of this vv-RKHS based approach is that one can benefit from the functional point of view, design new regularizers and impose various constraints on the whole continuum of tasks, e. g.,

- The continuity of the $\theta \mapsto \mathbf{h}(x)(\theta)$ function is a natural desirable property: for a given input x , the predictions on similar tasks should also be similar.
- Another example is to impose a shape constraint in QR: the conditional quantile should be increasing w. r. t. the hyperparameter θ . This requirement can be imposed through the functional view of the problem.
- In DLSE, to get nested level sets, one would want that for all $x \in \mathcal{X}$, the decision function $\theta \mapsto \mathbb{1}_{\mathbb{R}_+}(\mathbf{h}(x)(\theta) - \mathbf{t}(\theta))$ changes its sign only once.

To keep the presentation simple, in the sequel we are going to focus on ITL in the supervised setting; unsupervised tasks can be handled similarly.

Assume that \mathbf{h} belongs to some space $\mathcal{H} \subseteq \mathcal{F}(\mathcal{X}; \mathcal{F}(\Theta; \mathcal{Y}))$ and introduce an integrated loss func-

tion

$$V(\mathbf{y}, \mathbf{h}(\mathbf{x})) := \int_{\Theta} v(\theta, \mathbf{y}, \mathbf{h}(\mathbf{x})(\theta)) d\mu(\theta), \quad (5)$$

where the local loss $v: \Theta \times \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ denotes v_{θ} seen as a function of three variables including the hyperparameter and μ is a probability measure on Θ which encodes the importance of the prediction at different hyperparameter values. Without prior information and for compact Θ , one may consider μ to be uniform. The true risk reads then

$$R(\mathbf{h}) := \mathbf{E}_{\mathbf{X}, \mathbf{Y}} [V(\mathbf{Y}, \mathbf{h}(\mathbf{X}))]. \quad (6)$$

Intuitively, minimizing the expectation of the integral over θ in a rich enough space corresponds to searching for a pointwise minimizer $\mathbf{x} \mapsto \mathbf{h}^*(\mathbf{x})(\theta)$ of the parametrized tasks introduced in Eq. (1) with, for instance, the implicit space constraint that $\theta \mapsto \mathbf{h}^*(\mathbf{x})(\theta)$ is a continuous function for each input \mathbf{x} . We show in Proposition 2.1 that this is precisely the case in QR:

Proposition 2.1. *Let \mathbf{X}, \mathbf{Y} be two random variables respectively taking values in \mathcal{X} and \mathbb{R} , and $\mathbf{q}: \mathcal{X} \rightarrow \mathcal{F}([0, 1], \mathbb{R})$ the associated conditional quantile function. Let μ be a positive measure on $[0, 1]$ such that $\int_0^1 \mathbf{E} [v_{\theta}(\mathbf{Y}, \mathbf{q}(\mathbf{X})(\theta))] d\mu(\theta) < \infty$. Then $\forall \mathbf{h} \in \mathcal{F}(\mathcal{X}; \mathcal{F}([0, 1]; \mathbb{R}))$*

$$R(\mathbf{h}) - R(\mathbf{q}) \geq 0,$$

where R is the risk defined in Eq. (6).

Interestingly, the empirical counterpart of the true risk minimization can now be considered with a much richer family of penalty terms:

$$\min_{\mathbf{h} \in \mathcal{H}} R_s(\mathbf{h}) + \Omega(\mathbf{h}), \quad R_s(\mathbf{h}) := \frac{1}{n} \sum_{i=1}^n V(\mathbf{y}_i, \mathbf{h}(\mathbf{x}_i)). \quad (7)$$

Here, $\Omega(\mathbf{h})$ can be a weighted sum of various penalties

- imposed directly on $(\theta, \mathbf{x}) \mapsto \mathbf{h}(\mathbf{x})(\theta)$, or
- integrated constraints on either $\theta \mapsto \mathbf{h}(\mathbf{x})(\theta)$ or $\mathbf{x} \mapsto \mathbf{h}(\mathbf{x})(\theta)$ such as

$$\int_{\mathcal{X}} \Omega_1(\mathbf{h}(\mathbf{x})(\cdot)) d\mathbf{P}(\mathbf{x}) \text{ or } \int_{\Theta} \Omega_2(\mathbf{h}(\cdot)(\theta)) d\mu(\theta)$$

which allow the property enforced by Ω_1 or Ω_2 to hold pointwise on \mathcal{X} or Θ respectively.

It is worthwhile to see a concrete example before turning to the numerical solution (Section 3): in quantile regression, the monotonicity assumption of the $\theta \mapsto \mathbf{h}(\mathbf{x})(\theta)$ function can be encoded by choosing Ω_1 as

$$\Omega_1(f) = \lambda_{nc} \int_{\Theta} |-(\partial f)(\theta)|_+ d\mu(\theta).$$

Many different models (\mathcal{H}) could be applied to solve this problem. In our work we consider Reproducing Kernel Hilbert Spaces as they offer a simple and principled way to define regularizers by the appropriate choice of kernels and exhibit a significant flexibility.

3 Solving the problem in RKHSs

This section is dedicated to solving the ITL problem defined in Eq. (7). In Section 3.1 we focus on the objective (\tilde{V}). The applied vv-RKHS model family is detailed in Section 3.2 with various penalty examples followed by representer theorems, giving rise to computational tractability.

3.1 Sampled Empirical Risk

In practice solving Eq. (7) can be rather challenging due to the integral over θ . One might consider different numerical integration techniques to handle this issue. We focus here on Quasi Monte Carlo (QMC) methods as they allow (i) efficient optimization over vv-RKHSs which we will use for modelling \mathcal{H} (Proposition 3.1), and (ii) enable us to derive generalization guarantees (Proposition 4.1). Indeed, let

$$\tilde{V}(\mathbf{y}, \mathbf{h}(\mathbf{x})) := \sum_{j=1}^m w_j v(\theta_j, \mathbf{y}, \mathbf{h}(\mathbf{x})(\theta_j)) \quad (8)$$

be the QMC approximation of Eq. (5). Let $w_j = m^{-1} F^{-1}(\theta_j)$, and $(\theta_j)_{j=1}^m$ be a sequence with values in $[0, 1]^d$ such as the Sobol or Halton sequence where μ is assumed to be absolutely continuous w.r.t. the Lebesgue measure and F is the associated cdf. Using this notation and the training samples $\mathcal{S} = ((\mathbf{x}_i, \mathbf{y}_i))_{i=1}^n$, the empirical risk takes the form

$$\tilde{R}_s(\mathbf{h}) := \frac{1}{n} \sum_{i=1}^n \tilde{V}(\mathbf{y}_i, \mathbf{h}(\mathbf{x}_i)) \quad (9)$$

and the problem to solve is

$$\min_{\mathbf{h} \in \mathcal{H}} \tilde{R}_s(\mathbf{h}) + \Omega(\mathbf{h}). \quad (10)$$

3.2 Hypothesis class (\mathcal{H})

Recall that $\mathcal{H} \subseteq \mathcal{F}(\mathcal{X}; \mathcal{F}(\Theta; \mathcal{Y}))$, in other words $\mathbf{h}(\mathbf{x})$ is a $\Theta \mapsto \mathcal{Y}$ function for all $\mathbf{x} \in \mathcal{X}$. In this work we assume that the $\Theta \mapsto \mathcal{Y}$ mapping can be described by an RKHS $\mathcal{H}_{k_{\Theta}}$ associated to a $k_{\Theta}: \Theta \times \Theta \rightarrow \mathbb{R}$ scalar-valued kernel defined on the hyperparameters. Let $k_{\mathcal{X}}: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a scalar-valued kernel on the input space. The $\mathbf{x} \mapsto$ (hyperparameter \mapsto output)

relation, i. e. $h: \mathcal{X} \rightarrow \mathcal{H}_{k_\Theta}$ is then modelled by the Vector-Valued Reproducing Kernel Hilbert Space $\mathcal{H}_K = \overline{\text{span}} \{ K(\cdot, x)f \mid x \in \mathcal{X}, f \in \mathcal{H}_{k_\Theta} \}$, where the operator-valued kernel K is defined as $K(x, z) = k_X(x, z)I$, and $I = I_{\mathcal{H}_{k_\Theta}}$ is the identity operator on \mathcal{H}_{k_Θ} .

This so-called decomposable Operator-Valued Kernel has several benefits and gives rise to a function space with a well-known structure. One can consider elements $h \in \mathcal{H}_K$ as mappings from \mathcal{X} to \mathcal{H}_{k_Θ} , and also as functions from $(\mathcal{X} \times \Theta)$ to \mathbb{R} . It is indeed known that there is an isometry between \mathcal{H}_K and $\mathcal{H}_{k_X} \otimes \mathcal{H}_{k_\Theta}$, the RKHS associated to the product kernel $k_X \otimes k_\Theta$. The equivalence between these views allows a great flexibility and enables one to follow a functional point of view (to analyse statistical aspects) or to leverage the tensor product point of view (to design new kind of penalization schemes). Below we detail various regularizers before focusing on the representer theorems.

- **Ridge Penalty:** For QR and CSC, a natural regularization is the squared vv-RKHS norm

$$\Omega^{\text{RIDGE}}(h) = \frac{\lambda}{2} \|h\|_{\mathcal{H}_K}^2, \quad \lambda > 0. \quad (11)$$

This choice is amenable to excess risk analysis (see Proposition 4.1). It can be also seen as the counterpart of the classical (multi-task regularization term introduced by Sangnier et al. [SFd16], compatible with an infinite number of tasks. $\|\cdot\|_{\mathcal{H}_K}^2$ acts by constraining the solution to a ball of a finite radius within the vv-RKHS, whose shape is controlled by both k_X and k_Θ .

- **$L^{2,1}$ -penalty:** For DLSE, it is more adequate to apply an $L^{2,1}$ -RKHS mixed regularizer:

$$\Omega^{\text{DLSE}}(h) = \frac{1}{2} \int_{\Theta} \|h(\cdot)(\theta)\|_{\mathcal{H}_{k_X}}^2 d\mu(\theta) \quad (12)$$

which is an example of a Θ -integrated penalty. This Ω choice allows the preservation of the θ -property (see Fig. 2), i. e. that the proportion of the outliers is θ .

- **Shape Constraints:** Taking the example of QR it is advantageous to ensure the monotonicity of the estimated quantile function. Let $\partial_\Theta h$ denotes the derivative of $h(x)(\theta)$ with respect to θ . Then one should solve

$$\begin{aligned} & \arg \min_{h \in \mathcal{H}_K} \tilde{R}_S(h) + \Omega^{\text{RIDGE}}(h) \\ & \text{s. t. } \forall (x, \theta) \in \mathcal{X} \times \Theta, (\partial_\Theta h)(x)(\theta) \geq 0. \end{aligned}$$

However, the functional constraint prevents a tractable optimization scheme. To mitigate this bottleneck, we penalize if the derivative of h w. r. t. θ

is negative:

$$\Omega_{\text{nc}}(h) := \lambda_{\text{nc}} \int_{\mathcal{X}} \int_{\Theta} |-(\partial_\Theta h)(x)(\theta)|_+ d\mu(\theta) d\mathbf{P}(x). \quad (13)$$

When $\mathbf{P} := \mathbf{P}_X$ this penalization can rely on the same anchors and weights as the ones used to approximate the integrated loss function:

$$\tilde{\Omega}_{\text{nc}}(h) = \lambda_{\text{nc}} \sum_{i,j=1}^{n,m} w_j |-(\partial_X h)(x_i)(\theta_j)|_+. \quad (14)$$

Thus, one can modify the overall regularizer in QR to be

$$\Omega(h) := \Omega^{\text{RIDGE}}(h) + \tilde{\Omega}_{\text{nc}}(h). \quad (15)$$

3.3 Representer Theorems

Apart from the flexibility of regularizer design, the other advantage of applying vv-RKHS as hypothesis class is that it gives rise to finite-dimensional representation of the ITL solution under mild conditions. The representer theorem Proposition 3.1 applies to CSC when $\lambda_{\text{nc}} = 0$ and to QR when $\lambda_{\text{nc}} > 0$.

Proposition 3.1 (Representer). *Assume that for $\forall \theta \in \Theta, v_\theta$ is a proper lower semicontinuous convex function with respect to its second argument. Then*

$$\arg \min_{h \in \mathcal{H}_K} \tilde{R}_S(h) + \Omega(h), \quad \lambda > 0$$

with $\Omega(h)$ defined as in Eq. (15), has a unique solution h^ , and $\exists (\alpha_{ij})_{i,j=1}^{n,m}, (\beta_{ij})_{i,j=1}^{n,m} \in \mathbb{R}^{2nm}$ such that $\forall x \in \mathcal{X}$*

$$h^*(x) = \sum_{i=1}^n k_X(x, x_i) \left(\sum_{j=1}^m \alpha_{ij} k_\Theta(\cdot, \theta_j) + \beta_{ij} (\partial_2 k_\Theta)(\cdot, \theta_j) \right).$$

Sketch of the proof. *First, we prove that the function to minimize is coercive, convex, lower semicontinuous, hence it has a unique minimum. Then \mathcal{H}_K is decomposed into two orthogonal subspaces and we use the reproducing property to get the finite representation.*

For DLSE, we similarly get a representer theorem with the following modelling choice. Let $k_b: \Theta \times \Theta \rightarrow \mathbb{R}$ be a scalar-valued kernel (possibly different from k_Θ), \mathcal{H}_{k_b} the associated RKHS and $t \in \mathcal{H}_{k_b}$. Assume also that $\Theta \subseteq [\epsilon, 1]$ where $\epsilon > 0$.^{*} Then, learning a continuum of level sets boils down to the minimization problem

$$\arg \min_{h \in \mathcal{H}_K, t \in \mathcal{H}_{k_b}} \tilde{R}_S(h, t) + \tilde{\Omega}(h, t), \quad \lambda > 0, \quad (16)$$

^{*}We choose $\Theta \subseteq [\epsilon, 1]$, $\epsilon > 0$ rather than $\Theta \subseteq [0, 1]$ because the loss might not be integrable on $[0, 1]$.

where $\tilde{\Omega}(\mathbf{h}, \mathbf{t}) = \frac{1}{2} \sum_{j=1}^m w_j \|\mathbf{h}(\cdot)(\theta_j)\|_{\mathcal{H}_{k_x}}^2 + \frac{\lambda}{2} \|\mathbf{t}\|_{\mathcal{H}_{k_b}}^2$, $\tilde{\mathcal{R}}_S(\mathbf{h}, \mathbf{t}) = \frac{1}{n} \sum_{i,j=1}^{n,m} \frac{w_j}{\theta_j} (|\mathbf{t}(\theta_j) - \mathbf{h}(x_i)(\theta_j)|_+ - \mathbf{t}(\theta_j))$.

Proposition 3.2 (Representer). *Assume that k_Θ is bounded: $\sup_{\theta \in \Theta} k_\Theta(\theta, \theta) < +\infty$. Then the minimization problem described in Eq. (16) has a unique solution $(\mathbf{h}^*, \mathbf{t}^*)$ and there exist $(\alpha_{ij})_{i,j=1}^{n,m} \in \mathbb{R}^{n \times m}$ and $(\beta_j)_{j=1}^m \in \mathbb{R}^m$ such that for $\forall(x, \theta) \in \mathcal{X} \times [\epsilon, 1]$,*

$$\begin{aligned} \mathbf{h}^*(x)(\theta) &= \sum_{i,j=1}^{n,m} \alpha_{ij} k_x(x, x_i) k_\Theta(\theta, \theta_j), \\ \mathbf{t}^*(\theta) &= \sum_{j=1}^m \beta_j k_b(\theta, \theta_j). \end{aligned}$$

Sketch of the proof. *First we show that the infimum exists, and that it must be attained in some subspace of $\mathcal{H}_K \times \mathcal{H}_{k_b}$ over which the objective function is coercive. By the reproducing property, we get the claimed finite decomposition.*

Remarks:

- Models with bias: it can be advantageous to add a bias to the model, which is here a function of the hyperparameter θ : $\mathbf{h}(x)(\theta) = f(x)(\theta) + \mathbf{b}(\theta)$, $f \in \mathcal{H}_K$, $\mathbf{b} \in \mathcal{H}_{k_b}$, where $k_b : \Theta \times \Theta \rightarrow \mathbb{R}$ is a scalar-valued kernel. This can be the case for example if the kernel on the hyperparameters is the constant kernel, i.e. $k_\Theta(\theta, \theta') = 1$ ($\forall \theta, \theta' \in \Theta$), hence the model $f(x)(\theta)$ would not depend on θ . An analogous statement to Proposition 3.1 still holds for the biased model if one adds a regularization $\lambda_b \|\mathbf{b}\|_{\mathcal{H}_{k_b}}^2$, $\lambda_b > 0$ to the risk.
- Relation to JQR: In ∞ -QR, by choosing k_Θ to be the Gaussian kernel, $k_b(x, z) = \mathbb{1}_{\{x\}}(z)$, $\mu = \frac{1}{m} \sum_{j=1}^m \delta_{\theta_j}$, where δ_θ is the Dirac measure concentrated on θ , one gets back Sangnier et al. [SFd16]’s Joint Quantile Regression (JQR) framework as a special case of our approach. In contrast to the JQR, however, in ∞ -QR one can predict the quantile value at any $\theta \in (0, 1)$, even outside the $(\theta_j)_{j=1}^m$ used for learning.
- Relation to q-OCSVM: In DLSE, by choosing $k_\Theta(\theta, \theta') = 1$ (for all $\theta, \theta' \in \Theta$) to be the constant kernel, $k_b(\theta, \theta') = \mathbb{1}_{\{\theta\}}(\theta')$, $\mu = \frac{1}{m} \sum_{j=1}^m \delta_{\theta_j}$, our approach specializes to q-OCSVM [GLM13].
- Relation to Kadri et al. [Kad+16]: Note that Operator-Valued Kernels for functional outputs have also been used in [Kad+16], under the form of integral operators acting on L^2 spaces. Both kernels give rise to the same space of functions; the benefit of our approach being to provide an *exact* finite representation of the solution (see Proposition 3.1).

- Efficiency of the decomposable kernel: this kernel choice allows to rewrite the expansions in Propositions 3.1 and 3.2 as a Kronecker products and the complexity of the prediction of n' points for m' quantile becomes $\mathcal{O}(m'n n + n'n m)$ instead of $\mathcal{O}(m'n n')$.

4 Excess Risk Bounds

Below we provide a generalization error analysis to the solution of Eq. (10) for QR and CSC (with Ridge regularization and without shape constraints) by stability argument [BE02], extending the work of Audiffren et al. [AK13] to Infinite-Task Learning. The proposition instantiates the guarantee for the QMC scheme.

Proposition 4.1 (Generalization). *Let $\mathbf{h}^* \in \mathcal{H}_K$ be the solution of Eq. (10) for the QR or CSC problem with QMC approximation. Under mild conditions on the kernels k_x, k_Θ and $\mathbf{P}_{X,Y}$, one has*

$$\mathbf{R}(\mathbf{h}^*) \leq \tilde{\mathcal{R}}_S(\mathbf{h}^*) + \mathcal{O}_{\mathbf{P}_{X,Y}} \left(\frac{1}{\sqrt{\lambda n}} \right) + \mathcal{O} \left(\frac{\log(m)}{\sqrt{\lambda m}} \right). \quad (17)$$

Sketch of the proof. *The error resulting from sampling $\mathbf{P}_{X,Y}$ and the inexact integration is respectively bounded by β -stability [Kad+16] and QMC results.[†]*

(n, m) Trade-off: The proposition reveals the interplay between the two approximations, n (the number of training samples) and m (the number of locations taken in the integral approximation), and allows to identify the regime in $\lambda = \lambda(n, m)$ driving the excess risk to zero. Indeed by choosing $m = \sqrt{n}$ and discarding logarithmic factors for simplicity, $\lambda \gg n^{-1}$ is sufficient. The mild assumptions imposed are: boundedness on both kernels and the random variable Y , as well as some smoothness of the kernels.

5 Numerical Examples

In this section we provide numerical examples illustrating the efficiency of the proposed ITL approach.[‡] We used the following datasets in our experiments:

- Quantile Regression: we used (i) a sine synthetic benchmark [SFd16]: a sine curve at 1Hz modulated by a sine envelope at 1/3Hz and mean 1, distorted with a Gaussian noise of mean 0 and a linearly decreasing standard deviation from 1.2 at $x = 0$ to 0.2 at $x = 1.5$.

[†]The QMC approximation may involve the Sobol sequence with discrepancy $m^{-1} \log(m)^s$ ($s = \dim(\Theta)$).

[‡]The code is available at <https://bitbucket.org/RomainBrault/itl>.

Table 1: Quantile Regression on 20 UCI datasets. Reported: $100 \times$ value of the pinball loss, $100 \times$ crossing loss (smaller is better). p.-val.: outcome of the Mann-Whitney-Wilcoxon test of JQR vs. ∞ -QR and Independent vs. ∞ -QR. Boldface: significant values w. r. t. ∞ -QR.

DATASET	JQR				IND-QR				∞ -QR	
	(PINBALL)	P.-VAL.)	(CROSS	P.-VAL.)	(PINBALL)	P.-VAL.)	(CROSS	P.-VAL.)	PINBALL	CROSS
COBARORE	159 ± 24	$9 \cdot 10^{-01}$	0.1 ± 0.4	$6 \cdot 10^{-01}$	150 ± 21	$2 \cdot 10^{-01}$	0.3 ± 0.8	$7 \cdot 10^{-01}$	165 ± 36	2.0 ± 6.0
ENGEL	175 ± 555	$6 \cdot 10^{-01}$	0.0 ± 0.2	$1 \cdot 10^{+00}$	63 ± 53	$8 \cdot 10^{-01}$	4.0 ± 12.8	$8 \cdot 10^{-01}$	47 ± 6	0.0 ± 0.1
BOSTONHOUSING	49 ± 4	$8 \cdot 10^{-01}$	0.7 ± 0.7	$2 \cdot 10^{-01}$	49 ± 4	$8 \cdot 10^{-01}$	1.3 ± 1.2	$1 \cdot 10^{-05}$	49 ± 4	0.3 ± 0.5
CAUTION	88 ± 17	$6 \cdot 10^{-01}$	0.1 ± 0.2	$6 \cdot 10^{-01}$	89 ± 19	$4 \cdot 10^{-01}$	0.3 ± 0.4	$2 \cdot 10^{-04}$	85 ± 16	0.0 ± 0.1
FTCOLLINSNOW	154 ± 16	$8 \cdot 10^{-01}$	0.0 ± 0.0	$6 \cdot 10^{-01}$	155 ± 13	$9 \cdot 10^{-01}$	0.2 ± 0.9	$8 \cdot 10^{-01}$	156 ± 17	0.1 ± 0.6
HIGHWAY	103 ± 19	$4 \cdot 10^{-01}$	0.8 ± 1.4	$2 \cdot 10^{-02}$	99 ± 20	$9 \cdot 10^{-01}$	6.2 ± 4.1	$1 \cdot 10^{-07}$	105 ± 36	0.1 ± 0.4
HEIGHTS	127 ± 3	$1 \cdot 10^{+00}$	0.0 ± 0.0	$1 \cdot 10^{+00}$	127 ± 3	$9 \cdot 10^{-01}$	0.0 ± 0.0	$1 \cdot 10^{+00}$	127 ± 3	0.0 ± 0.0
SNIFFER	43 ± 6	$8 \cdot 10^{-01}$	0.1 ± 0.3	$2 \cdot 10^{-01}$	44 ± 5	$7 \cdot 10^{-01}$	1.4 ± 1.2	$6 \cdot 10^{-07}$	44 ± 7	0.1 ± 0.1
SNOWGEESE	55 ± 20	$7 \cdot 10^{-01}$	0.3 ± 0.8	$3 \cdot 10^{-01}$	53 ± 18	$6 \cdot 10^{-01}$	0.4 ± 1.0	$5 \cdot 10^{-02}$	57 ± 20	0.2 ± 0.6
UFC	81 ± 5	$6 \cdot 10^{-01}$	0.0 ± 0.0	$4 \cdot 10^{-04}$	82 ± 5	$7 \cdot 10^{-01}$	1.0 ± 1.4	$2 \cdot 10^{-04}$	82 ± 4	0.1 ± 0.3
BigMAC2003	80 ± 21	$7 \cdot 10^{-01}$	1.4 ± 2.1	$4 \cdot 10^{-04}$	74 ± 24	$9 \cdot 10^{-02}$	0.9 ± 1.1	$7 \cdot 10^{-05}$	84 ± 24	0.2 ± 0.4
UN3	98 ± 9	$8 \cdot 10^{-01}$	0.0 ± 0.0	$1 \cdot 10^{-01}$	99 ± 9	$1 \cdot 10^{+00}$	1.2 ± 1.0	$1 \cdot 10^{-05}$	99 ± 10	0.1 ± 0.4
BIRTHWT	141 ± 13	$1 \cdot 10^{+00}$	0.0 ± 0.0	$6 \cdot 10^{-01}$	140 ± 12	$9 \cdot 10^{-01}$	0.1 ± 0.2	$7 \cdot 10^{-02}$	141 ± 12	0.0 ± 0.0
CRABS	11 ± 1	$4 \cdot 10^{-05}$	0.0 ± 0.0	$8 \cdot 10^{-01}$	11 ± 1	$2 \cdot 10^{-04}$	0.0 ± 0.0	$2 \cdot 10^{-05}$	13 ± 3	0.0 ± 0.0
GAGURINE	61 ± 7	$4 \cdot 10^{-01}$	0.0 ± 0.1	$3 \cdot 10^{-03}$	62 ± 7	$5 \cdot 10^{-01}$	0.1 ± 0.2	$4 \cdot 10^{-04}$	62 ± 7	0.0 ± 0.0
GEYSER	105 ± 7	$9 \cdot 10^{-01}$	0.1 ± 0.3	$9 \cdot 10^{-01}$	105 ± 6	$9 \cdot 10^{-01}$	0.2 ± 0.3	$6 \cdot 10^{-01}$	104 ± 6	0.1 ± 0.2
GILGAIS	51 ± 6	$5 \cdot 10^{-01}$	0.1 ± 0.1	$1 \cdot 10^{-01}$	49 ± 6	$6 \cdot 10^{-01}$	1.1 ± 0.7	$2 \cdot 10^{-05}$	49 ± 7	0.3 ± 0.3
TOPO	69 ± 18	$1 \cdot 10^{+00}$	0.1 ± 0.5	$1 \cdot 10^{+00}$	71 ± 20	$1 \cdot 10^{+00}$	1.7 ± 1.4	$3 \cdot 10^{-07}$	70 ± 17	0.0 ± 0.0
MCYCLE	66 ± 9	$9 \cdot 10^{-01}$	0.2 ± 0.3	$7 \cdot 10^{-03}$	66 ± 8	$9 \cdot 10^{-01}$	0.3 ± 0.3	$7 \cdot 10^{-06}$	65 ± 9	0.0 ± 0.1
CPUS	7 ± 4	$2 \cdot 10^{-04}$	0.7 ± 1.0	$5 \cdot 10^{-04}$	7 ± 5	$3 \cdot 10^{-04}$	1.2 ± 0.8	$6 \cdot 10^{-08}$	16 ± 10	0.0 ± 0.0




Figure 1: Impact of crossing penalty on toy data. Left plot: strong non-crossing penalty ($\lambda_{nc} = 10$). Right plot: no non-crossing penalty ($\lambda_{nc} = 0$). The plots show 100 quantiles of the continuum learned, linearly spaced between 0 (blue) and 1 (red). Notice that the non-crossing penalty does not provide crossings to occur in the regions where there is no points to enforce the penalty (e.g. $x \in [0.13, 0.35]$). This phenomenon is alleviated by the regularity of the model.

(ii) 20 standard regression datasets from UCI. The number of samples varied between 38 (CobarOre) and 1375 (Height). The observations were standardised to have unit variance and zero mean for each attribute.

- **Density Level-Set Estimation:** The Wilt database from the UCI repository with 4839 samples and 5 attributes, and the Spambase UCI dataset with 4601 samples and 57 attributes served as benchmarks.

Note on Optimization: There are several ways to solve the non-smooth optimization problems associated to the QR, DLSE and CSC tasks. One could proceed for example by duality—as it was done in JQR Sangnier et al. [SFd16]—, or apply sub-gradient descent techniques (which often converge quite slowly). In order to allow unified treatment and efficient solution in our experiments we used the L-BFGS-B [Zhu+97] optimization scheme which is widely popular in large-scale learning, with non-smooth extensions [KW17]. The technique requires only evaluation of objective function along with its gradient, which can be computed automatically using reverse mode automatic differentiation (as in Abadi et al. [Aba+16]). To benefit from the available fast smooth implementations [JOP+01], we applied an infimal convolution on the non-differentiable terms of the objective. Under the assumption that $m = \mathcal{O}(\sqrt{n})$ (see Proposition 4.1), the complexity per L-BFGS-B iteration is $\mathcal{O}(n^2\sqrt{n})$.

QR: The efficiency of the non-crossing penalty is illustrated in Fig. 1 on the synthetic sine wave dataset described in Section 5 where $n = 40$ and $m = 20$ points have been generated. Many crossings are visible on the right plot, while they are almost not noticeable on the left plot, using the non-crossing penalty. Concerning our real-world examples, to study the efficiency of the proposed scheme in quantile regression the following experimental protocol was applied. Each dataset (Section 5) was splitted randomly into a training set (70%) and a test set (30%). We optimized the hyperparameters by minimizing a 5-folds cross validation with a Bayesian optimizer[§]. Once the hyperparameters were obtained, a new regressor was learned on the whole training set using the optimized hyperparameters. We report the value of the pinball loss and the crossing loss on the test set for three methods: our technique is called ∞ -QR, we refer to Sangnier et al. [SFd16]’s approach as JQR, and independent learning (abbreviated as IND-QR) represents a further baseline.

We repeated 20 simulations (different random training-test splits); the results are also compared using

a Mann-Whitney-Wilcoxon test. A summary is provided in Table 1. Notice that while JQR is tailored to predict finite many quantiles, our ∞ -QR method estimates the *whole quantile function* hence solves a more challenging task. Despite the more difficult problem solved, as Table 1 suggest that the performance in terms of pinball loss of ∞ -QR is comparable to that of the state-of-the-art JQR on all the twenty studied benchmarks, except for the ‘crabs’ and ‘cpus’ datasets (p-val. $< 0.25\%$). In addition, when considering the non-crossing penalty one can observe that ∞ -QR outperforms the IND-QR baseline on eleven datasets (p-val. $< 0.25\%$) and JQR on two datasets. This illustrates the efficiency of the constraint based on the continuum scheme.

CSC: To illustrate the advantage of (infinite) joint learning we used two synthetic datasets CIRCLES and TWO-MOONS and the UCI IRIS dataset. We chose k_x to be a Gaussian kernel with bandwidth $\sigma_x = (2\gamma_x)^{(-1/2)}$ the median of the Euclidean pairwise distances of the input points [JDH99]. k_θ is also a Gaussian kernel with bandwidth $\gamma_\theta = 5$. We used $m = 20$ for all datasets. As a baseline we trained independently 3 Cost-Sensitive Classification classifiers with $\theta \in \{-0.9, 0, 0.9\}$. We repeated 50 times a random 50 – 50% train-test split of the dataset and report the average test error and standard deviation (in terms of sensitivity and specificity)

Our results are illustrated in Table 2. For $\theta = -0.9$, both independent and joint learners give the desired 100% specificity; the joint Cost-Sensitive Classification scheme however has significantly higher sensitivity value (15% vs 0%) on the dataset CIRCLES. Similar conclusion holds for the $\theta = +0.9$ extreme: the ideal sensitivity is reached by both techniques, but the joint learning scheme performs better in terms of specificity (0% vs 12%) on the dataset CIRCLES.

DLSE: To assess the quality of the estimated model by ∞ -OCSVM, we illustrate the θ -property [Sch+00]: the proportion of inliers has to be approximately $1 - \theta$ ($\forall \theta \in (0, 1)$). For the studied datasets (Wilt, Spambase) we used the raw inputs without applying any preprocessing. Our input kernel was the exponentiated χ^2 kernel $k_x(x, z) := \exp\left(-\gamma_x \sum_{k=1}^d (x_k - z_k)^2 / (x_k + z_k)\right)$ with bandwidth $\gamma_x = 0.25$. A Gauss-Legendre quadrature rule provided the integral approximation in Eq. (8),

[§]We used a Gaussian Process model and minimized the Expected improvement. The optimizer was initialized using 27 samples from a Sobol sequence and ran for 50 iterations.

Table 2: ∞ -CSC vs Independent (IND)-CSC. Higher is better.

DATASET	METHOD	$\theta = -0.9$		$\theta = 0$		$\theta = +0.9$	
		SENSITIVITY	SPECIFICITY	SENSITIVITY	SPECIFICITY	SENSITIVITY	SPECIFICITY
TWO-MOONS	IND	0.3 \pm 0.05	0.99 \pm 0.01	0.83 \pm 0.03	0.86 \pm 0.03	0.99 \pm 0	0.32 \pm 0.06
	∞ -CSC	0.32 \pm 0.05	0.99 \pm 0.01	0.84 \pm 0.03	0.87 \pm 0.03	1 \pm 0	0.36 \pm 0.04
CIRCLES	IND	0 \pm 0	1 \pm 0	0.82 \pm 0.02	0.84 \pm 0.03	1 \pm 0	0 \pm 0
	∞ -CSC	0.15 \pm 0.05	1 \pm 0	0.82 \pm 0.02	0.84 \pm 0.03	1 \pm 0	0.12 \pm 0.05
IRIS	IND	0.88 \pm 0.08	0.94 \pm 0.06	0.94 \pm 0.05	0.92 \pm 0.06	0.97 \pm 0.05	0.87 \pm 0.06
	∞ -CSC	0.89 \pm 0.08	0.94 \pm 0.05	0.94 \pm 0.06	0.92 \pm 0.05	0.97 \pm 0.04	0.90 \pm 0.05
TOY	IND	0.51 \pm 0.06	0.98 \pm 0.01	0.83 \pm 0.03	0.86 \pm 0.03	0.97 \pm 0.01	0.49 \pm 0.07
	∞ -CSC	0.63 \pm 0.04	0.96 \pm 0.01	0.83 \pm 0.03	0.85 \pm 0.03	0.95 \pm 0.02	0.61 \pm 0.04

with $m = 100$ samples. We chose the Gaussian kernel for k_Θ ; its bandwidth parameter γ_Θ was the 0.2-quantile of the pairwise Euclidean distances between the θ_j 's obtained via the quadrature rule. The margin (bias) kernel was $k_b = k_\Theta$. As it can be seen in Fig. 2, the θ -property holds for the estimate which illustrates the efficiency of the proposed continuum approach for density level-set estimation.




Figure 2: Density Level-Set Estimation: the θ -property is approximately satisfied.

6 Conclusion

In this work we proposed Infinite Task Learning, a novel nonparametric framework aiming at jointly solving parametrized tasks for a continuum of hyperparameters. We provided excess risk guarantees for the studied ITL scheme, and demonstrated its practical efficiency and flexibility in various tasks including cost-sensitive classification, quantile regression and density level set estimation.

Acknowledgments

The authors thank Arthur Tenenhaus for some insightful discussions. This work was supported by

the Labex *DigiCosme* (project ANR-11-LABEX-0045-DIGICOSME) and the industrial chair *Machine Learning for Big Data* at Télécom ParisTech.

References

- [Aba+16] M. Abadi et al. “Tensorflow: Large-scale machine learning on heterogeneous distributed systems”. In: *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 2016, pp. 265–283 (cit. on p. 8).
- [AK13] J. Audiffren and H. Kadri. “Stability of Multi-Task Kernel Regression Algorithms”. In: *Asian Conference on Machine Learning (ACML)*. Vol. 29. PMLR, 2013, pp. 1–16 (cit. on p. 6).
- [ÁRL12] M. A. Álvarez, L. Rosasco, and N. D. Lawrence. “Kernels for vector-valued functions: a review”. In: *Foundations and Trends in Machine Learning* 4.3 (2012), pp. 195–266 (cit. on p. 3).
- [Bac+12] F. Bach et al. “Optimization with sparsity-inducing penalties”. In: *Foundations and Trends in Machine Learning* 4.1 (2012), pp. 1–106 (cit. on p. 1).
- [BE02] O. Bousquet and A. Elisseeff. “Stability and generalization”. In: *Journal of Machine Learning Research* 2 (2002), pp. 499–526 (cit. on p. 6).
- [Car+10] C. Carmeli et al. “Vector valued reproducing kernel Hilbert spaces and universality”. In: *Analysis and Applications* 8 (1 2010), pp. 19–61 (cit. on p. 2).

- [Cil+17] C. Ciliberto et al. “Consistent multitask learning with nonlinear output relations”. In: *Advances in Neural Information Processing Systems (NIPS)*. 2017, pp. 1986–1996 (cit. on p. 3).
- [EMP05] T. Evgeniou, C. A. Micchelli, and M. Pontil. “Learning Multiple Tasks with kernel methods”. In: *JMLR* 6 (2005), pp. 615–637 (cit. on p. 3).
- [EP04] T. Evgeniou and M. Pontil. “Regularized multi-task learning”. In: *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM. 2004, pp. 109–117 (cit. on p. 1).
- [GLM13] A. Glazer, M. Lindenbaum, and S. Markovitch. “q-OCSVM: A q-quantile estimator for high-dimensional distributions”. In: *Advances in Neural Information Processing Systems (NIPS)*. 2013, pp. 503–511 (cit. on pp. 1–3, 6).
- [Has+04] T. Hastie et al. “The entire regularization path for the support vector machine”. In: *Journal of Machine Learning Research* 5.Oct (2004), pp. 1391–1415 (cit. on p. 1).
- [JDH99] T. Jaakkola, M. Diekhans, and D. Hausler. “Using the Fisher kernel method to detect remote protein homologies.” In: *ISMB*. Vol. 99. 1999, pp. 149–158 (cit. on p. 8).
- [JOP+01] E. Jones, T. Oliphant, P. Peterson, et al. *SciPy: Open source scientific tools for Python*. 2001 (cit. on p. 8).
- [Kad+16] H. Kadri et al. “Operator-valued Kernels for Learning from Functional Response Data”. In: *Journal of Machine Learning Research* 17 (2016), pp. 1–54 (cit. on p. 6).
- [KB78] R. Koenker and G. Bassett Jr. “Regression quantiles”. In: *Econometrica: journal of the Econometric Society* (1978), pp. 33–50 (cit. on pp. 1, 2).
- [KW17] N. Keskar and A. Wächter. “A limited-memory quasi-Newton algorithm for bound-constrained non-smooth optimization”. In: *Optimization Methods and Software* (2017), pp. 1–22 (cit. on p. 8).
- [Ped57] G. Pedrick. “Theory of reproducing kernels for Hilbert spaces of vector-valued functions”. PhD thesis. University of Kansas, 1957 (cit. on p. 2).
- [Sch+00] B. Schölkopf et al. “New support vector algorithms”. In: *Neural computation* 12.5 (2000), pp. 1207–1245 (cit. on pp. 1, 3, 8).
- [SFd16] M. Sangnier, O. Fercoq, and F. d’Alché-Buc. “Joint quantile regression in vector-valued RKHSs”. In: *Advances in Neural Information Processing Systems (NIPS)* (2016), pp. 3693–3701 (cit. on pp. 2, 3, 5, 6, 8).
- [TA77] A. N. Tikhonov and V. Y. Arsenin. *Solution of Ill-posed Problems*. Winston & Sons, 1977 (cit. on p. 1).
- [Tak+06] I. Takeuchi et al. “Nonparametric quantile estimation”. In: *Journal of Machine Learning Research* 7 (2006), pp. 1231–1264 (cit. on p. 1).
- [Tak+13] I. Takeuchi et al. “Parametric task learning”. In: *Advances in Neural Information Processing Systems (NIPS)*. 2013, pp. 1358–1366 (cit. on pp. 1, 3).
- [Tib96] R. Tibshirani. “Regression shrinkage and selection via the Lasso”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* (1996), pp. 267–288 (cit. on p. 1).
- [VGS97] V. Vapnik, S. E. Golowich, and A. J. Smola. “Support vector method for function approximation, regression estimation and signal processing”. In: *Advances in Neural Information Processing Systems (NIPS)*. 1997, pp. 281–287 (cit. on p. 1).
- [VV06] R. Vert and J.-P. Vert. “Consistency and convergence rates of one-class SVMs and related algorithms”. In: *Journal of Machine Learning Research* 7 (2006), pp. 817–854 (cit. on p. 3).
- [ZE01] B. Zadrozny and C. Elkan. “Learning and making decisions when costs and probabilities are both unknown”. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. 2001, pp. 204–213 (cit. on p. 1).
- [Zhu+97] C. Zhu et al. “Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization”. In: *ACM Transactions on Mathematical Software (TOMS)* 23.4 (1997), pp. 550–560 (cit. on p. 8).