

---

# The Advantage of Conditional Meta-Learning for Biased Regularization and Fine Tuning

---

Giulia Denevi<sup>1,\*</sup> Massimiliano Pontil<sup>1,2</sup> Carlo Ciliberto<sup>3</sup>

<sup>1</sup> University College of London (UK)

<sup>2</sup> Istituto Italiano di Tecnologia (Italy)

<sup>3</sup> Imperial College of London (UK)

\* Work done while the first author was with Istituto Italiano di Tecnologia (Italy)  
*g.denevi@ucl.ac.uk, massimiliano.pontil@iit.it, c.ciliberto@imperial.ac.uk*

## Abstract

Biased regularization and fine tuning are two recent meta-learning approaches. They have been shown to be effective to tackle distributions of tasks, in which the tasks' target vectors are all close to a common meta-parameter vector. However, these methods may perform poorly on heterogeneous environments of tasks, where the complexity of the tasks' distribution cannot be captured by a single meta-parameter vector. We address this limitation by conditional meta-learning, inferring a conditioning function mapping task's side information into a meta-parameter vector that is appropriate for that task at hand. We characterize properties of the environment under which the conditional approach brings a substantial advantage over standard meta-learning and we highlight examples of environments, such as those with multiple clusters, satisfying these properties. We then propose a convex meta-algorithm providing a comparable advantage also in practice. Numerical experiments confirm our theoretical findings.

## 1 Introduction

Biased regularization and fine tuning [5, 13, 14, 16, 17, 19, 23, 24, 28, 31, 33] are two recent optimization-based meta-learning techniques that transfer knowledge across an environment of tasks by leveraging a common meta-parameter vector. Their origin and inspiration go back to multi-task and transfer learning methods [10, 15, 27], designed to address a prescribed set of tasks with low variance. These techniques can be described as a nested optimization scheme: while at the within-task level, an inner algorithm performs tasks' specific optimization with the current meta-parameter vector, at the meta-level a meta-algorithm updates the aforementioned meta-parameter by leveraging the experience accumulated from the tasks observed so far. In biased regularization the inner algorithm is given by the within-task regularized empirical risk minimizer and the meta-parameter vector plays the role of a bias in the regularizer, while fine tuning employs online gradient descent as the within-task algorithm and the meta-parameter vector is the associated starting point.

Despite their success, the above methods may fail to adapt to heterogenous environments of tasks, in which the complexity of the tasks' distribution cannot be captured by a single meta-parameter vector. In literature, a variety of methods have tried to address this limitation by clustering the tasks and, then, leveraging tasks' similarities within each cluster [2, 4, 21, 29, 30]. However, such methods usually lead to non-convex formulations [2, 4] or provide only partial guarantees on surrogate convex problems [21, 30]. As alternative, recent approaches in meta-learning literature advocated learning a conditioning function that maps a task's dataset into a meta-parameter vector that is appropriate for the task at hand [9, 22, 35, 39–41]. This perspective has been shown to be promising in applications, however theoretical investigations are still lacking. The authors of [40] have recently made a first

step towards this theoretical investigation adopting a structured prediction perspective for conditional meta-learning. However, their method relies on a non-convex formulation of the problem. This allows them to provide guarantees only on the generalization properties of their method, but not on its estimation properties. In this work, we address the limitation above for biased regularization and fine tuning by developing a new conditional meta-learning framework. Specifically, we consider an environment of tasks provided with additional side information and we learn a conditioning function mapping task’s side information into a task’s specific meta-parameter vector. We then provide a statistical analysis demonstrating the potential advantage of our method over standard meta-learning.

**Contributions and organization.** Our work offers four contributions. First, in [Sec. 2](#), we introduce a new conditional meta-learning framework with side information for biased regularization and fine tuning. Second, in [Sec. 3](#), we formally show that, under certain assumptions, this conditional meta-learning approach results to be significantly advantageous w.r.t. the standard unconditional counterpart. We then describe two common settings in which such conditions are satisfied, supporting the potential importance of our study for real-world scenarios. Third, in [Sec. 4](#), we propose a convex meta-algorithm providing a comparable advantage also in practice, as the number of observed tasks increases. Fourth, in [Sec. 5](#), we present numerical experiments in which we test our theory and the performance of our method. Our conclusions are drawn in [Sec. 6](#) and technical proofs are postponed to the appendix.

## 2 Conditional meta-learning

In this section we describe and contrast the conditional meta-learning setting with side information to standard meta-learning. We first introduce the class of inner algorithms we consider in this work.

**Inner algorithms (linear supervised learning).** Let  $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$  with  $\mathcal{X} \subseteq \mathbb{R}^d$  and  $\mathcal{Y} \subseteq \mathbb{R}$  input and output spaces, respectively. Let  $\mathcal{P}(\mathcal{Z})$  be the set of probability distributions (tasks) over  $\mathcal{Z}$ . Given  $\mu \in \mathcal{P}(\mathcal{Z})$  and a loss function  $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ , our goal is to find a weight vector  $w_\mu \in \mathbb{R}^d$  minimizing the *expected risk*

$$\min_{w \in \mathbb{R}^d} \mathcal{R}_\mu(w) \quad \mathcal{R}_\mu(w) = \mathbb{E}_{(x,y) \sim \mu} \ell(\langle x, w \rangle, y), \quad (1)$$

where,  $\langle x, w \rangle$  denotes the standard inner product between  $x$  and  $w \in \mathbb{R}^d$ . In practice,  $\mu$  is unknown and only accessible through a training dataset  $Z = (x_i, y_i)_{i=1}^n \sim \mu^n$  of i.i.d. (identically independently distributed) points sampled from  $Z$ . The goal of a learning algorithm is to find a candidate weight vector incurring a small expected risk converging to the ideal  $\mathcal{R}_\mu(w_\mu)$  as  $n$  grows.

In this work we will focus on the family of learning algorithms performing biased regularized empirical risk minimization. Formally, given  $\mathcal{D} = \bigcup_{n \in \mathbb{N}} \mathcal{Z}^n$  the space of all datasets (of any finite cardinality  $n$ ) on  $\mathcal{Z}$  and a bias vector  $\theta \in \Theta = \mathbb{R}^d$ , we will consider learning algorithms  $A(\theta, \cdot) : \mathcal{D} \rightarrow \mathbb{R}^d$  such that,

$$A(\theta, Z) = \operatorname{argmin}_{w \in \mathbb{R}^d} \mathcal{R}_{Z,\theta}(w) \quad \mathcal{R}_{Z,\theta}(w) = \frac{1}{n} \sum_{i=1}^n \ell(\langle x_i, w \rangle, y_i) + \frac{\lambda}{2} \|w - \theta\|^2, \quad (2)$$

for any  $Z = (x_i, y_i)_{i=1}^n$ . Here  $\|\cdot\|$  denotes the Euclidean norm on  $\mathbb{R}^d$  and  $\lambda > 0$  is a regularization parameter encouraging the algorithm  $A(\theta, \cdot)$  to predict weight vectors that are close to  $\theta$ . We denote by  $\mathcal{R}_Z(\cdot) = 1/n \sum_{i=1}^n \ell(\langle x_i, w \rangle, y_i)$  the empirical risk associated to  $Z$ .

**Remark 1** (Fine tuning). *In this work we primarily focus on the family (2) of batch inner algorithms. However, following [13, 14], it is possible to extend our analysis to fine-tuning algorithms performing online gradient descent on  $\mathcal{R}_{Z,\theta}$ , with starting point  $w_1 = \theta \in \mathbb{R}^d$ , namely*

$$A(\theta, Z) = \frac{1}{n} \sum_{i=1}^n w_i, \quad w_{i+1} = w_i - \frac{s_i x_i - \lambda(w_i - \theta)}{\lambda_i}, \quad s_i \in \partial \ell(\cdot, y_i)(\langle x_i, w_i \rangle). \quad (3)$$

**(Unconditional) meta-learning.** Given a meta-distribution  $\rho \in \mathcal{P}(\mathcal{M})$  (or *environment* [7]) over a family  $\mathcal{M} \subseteq \mathcal{P}(\mathcal{Z})$  of distributions (tasks)  $\mu$ , meta-learning aims to learn an inner algorithm in the family that is well suited to tasks  $\mu$  sampled from  $\rho$ . This goal can be reformulated as finding a meta-parameter  $\theta_\rho \in \Theta$  whose associated algorithm  $A(\theta_\rho, \cdot)$  minimizes the *transfer risk*

$$\min_{\theta \in \Theta} \mathcal{E}_\rho(\theta) \quad \mathcal{E}_\rho(\theta) = \mathbb{E}_{\mu \sim \rho} \mathbb{E}_{Z \sim \mu^n} \mathcal{R}_\mu(A(\theta, Z)). \quad (4)$$

Standard meta-learning methods [5, 13, 14, 17–19, 24] usually address this problem via stochastic methods. They iteratively sample a task  $\mu \sim \rho$  and a dataset  $Z \sim \mu^n$ , and, then, they perform a step of stochastic gradient descent on a surrogate problem of (4) computed by using  $Z$ .

Although remarkably effective in many applications [5, 13, 14, 17–19, 24], the framework above implicitly assumes that a single bias vector is sufficient for the entire family of tasks sampled from  $\rho$ . Since this assumption may not hold for more complex meta-distributions (e.g. multi-clusters), recent works have advocated a conditional perspective to tackle this problem [9, 22, 35, 39–41].

**Conditional meta-learning.** Assume now that when sampling a task  $\mu$ , we are also given additional side information  $s \in \mathcal{S}$  to help solving the task. Within this setting the environment corresponds to a distribution  $\rho \in \mathcal{P}(\mathcal{M}, \mathcal{S})$  over the set  $\mathcal{M}$  of tasks and the set  $\mathcal{S}$  of possible side information. The notion of side information is general, and recovers settings where  $s$  contains descriptive features associated to a task (e.g. attributes in collaborative filtering [1]), additional information about the users in recommendation systems [20], or  $s$  is an additional dataset sampled from  $\mu$  (see [40] or [Rem. 2](#) below). Intuitively, meta-learning might solve a new task better if it was able to leverage this additional side information. We formalize this concept by adapting (or conditioning) the meta-parameters  $\theta \in \Theta$  on the side information  $s \in \mathcal{S}$ , by learning a meta-parameter-valued function  $\tau$  minimizing

$$\min_{\tau \in \mathcal{T}} \mathcal{E}_\rho(\tau) \quad \mathcal{E}_\rho(\tau) = \mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^n} \mathcal{R}_\mu(A(\tau(s), Z)), \quad (5)$$

over the space  $\mathcal{T}$  of measurable functions  $\tau : \mathcal{S} \rightarrow \Theta$ . Note that the unconditional meta-learning problem in (4) is retrieved by restricting (5) to  $\mathcal{T}^{\text{const}} = \{\tau \mid \tau(\cdot) \equiv \theta, \theta \in \Theta\}$ , the set of constant functions associating any side information to a fixed bias vector. We assume  $\rho$  to decompose in  $\rho(\cdot|s)\rho_{\mathcal{S}}(\cdot)$  and  $\rho(\cdot|\mu)\rho_{\mathcal{M}}(\cdot)$  the conditional and marginal distributions w.r.t. (with respect to)  $\mathcal{S}$  and  $\mathcal{M}$ . In the following, we will quantify the benefits of adopting the conditional perspective above and, then, we propose an efficient algorithm to address (5). We conclude this section by drawing a connection between our formulation and previous work on the topic.

**Remark 2** (Datasets as side information). *A relevant setting is the case where the side information  $s$  corresponds to an additional (conditional) dataset  $Z^{\text{cond}}$  sampled from  $\mu$ , as proposed in [40]. We note however that our sampling scheme in (5) implies that side information  $s$  and training set  $Z$  are independent conditioned on  $\mu$ . Hence, our framework does not allow having  $s = Z^{\text{cond}} = Z$ , namely, to use the same dataset for both conditioning and training the inner algorithm  $A(\tau(Z), Z)$ , as done in [40]. This is a minor issue since one can always split  $Z$  in two parts and use one part for training and the other one for conditioning.*

### 3 The advantage of conditional meta-learning

In this section we study the generalization properties of a given conditional function  $\tau$ . This will allow us to characterize the behavior of the ideal solution of (5) and to illustrate the potential advantage of conditional meta-learning. Specifically, we wish to estimate the error  $\mathcal{E}_\rho(\tau)$  w.r.t. the ideal risk

$$\mathcal{E}_\rho^* = \mathbb{E}_{\mu \sim \rho} \mathcal{R}_\mu(w_\mu) \quad w_\mu = \operatorname{argmin}_{w \in \mathbb{R}^d} \mathcal{R}_\mu(w). \quad (6)$$

For any  $\tau \in \mathcal{T}$  the following quantity will play a central role in our analysis:

$$\operatorname{Var}_\rho(\tau)^2 = \mathbb{E}_{(\mu, s) \sim \rho} \|w_\mu - \tau(s)\|^2. \quad (7)$$

With some abuse of terminology, we refer to  $\operatorname{Var}_\rho(\tau)$  as the *variance* of  $w_\mu$  w.r.t.  $\tau$  (it corresponds to the actual variance of  $w_\mu$  when  $\tau$  is the minimizer, see [Lemma 2](#) below). Under the following assumption, we can control the excess risk of  $\tau$  in terms of  $\operatorname{Var}_\rho(\tau)$ .

**Assumption 1.** *Let  $\ell$  be a convex and  $L$ -Lipschitz loss function in the first argument. Additionally, there exist  $R > 0$  such that  $\|x\| \leq R$  for any  $x \in \mathcal{X}$ .*

**Theorem 1** (Excess risk with generic conditioning function  $\tau$ ). *Let [Asm. 1](#) hold. Given  $\tau \in \mathcal{T}$ , let  $A(\theta, \cdot)$  be the generic inner algorithm in (2) with regularization parameter  $\lambda = 2LR\operatorname{Var}_\rho(\tau)^{-1}n^{-1/2}$ . Then,*

$$\mathcal{E}_\rho(\tau) - \mathcal{E}_\rho^* \leq \frac{2RL \operatorname{Var}_\rho(\tau)}{n^{1/2}}. \quad (8)$$

**Proof.** We consider the decomposition  $\mathcal{E}_\rho(\tau) - \mathcal{E}_\rho^* = \mathbb{E}_{(\mu,s) \sim \rho} [\mathbf{B}_{\mu,s} + \mathbf{C}_{\mu,s}]$ , with

$$\mathbf{B}_{\mu,s} = \mathbb{E}_{Z \sim \mu^n} [\mathcal{R}_\mu(A(\tau(s), Z)) - \mathcal{R}_Z(A(\tau(s), Z))] \quad (9)$$

$$\mathbf{C}_{\mu,s} = \mathbb{E}_{Z \sim \mu^n} [\mathcal{R}_Z(A(\tau(s), Z)) - \mathcal{R}_\mu(w_\mu)] \leq \mathbb{E}_{Z \sim \mu^n} \left[ \min_{w \in \mathbb{R}^d} \mathcal{R}_{Z, \tau(s)}(w) - \mathcal{R}_\mu(w_\mu) \right]. \quad (10)$$

$\mathbf{B}_{\mu,s}$  is the generalization error of the inner algorithm  $A(\tau(s), \cdot)$  on the task  $\mu$ . Hence, applying [Asm. 1](#) and the stability arguments in [Prop. 5](#) in [App. A](#), we can write  $\mathbf{B}_{\mu,s} \leq 2R^2L^2(\lambda n)^{-1}$ . Regarding the term  $\mathbf{C}_{\mu,s}$ , exploiting the definition of the algorithm in [\(2\)](#), we can write  $\mathbf{C}_{\mu,s} \leq \frac{\lambda}{2} \|w_\mu - \tau(s)\|^2$ . The desired statement follows by combining the two bounds above and optimizing w.r.t.  $\lambda$ .  $\blacksquare$

[Thm. 1](#) suggests that a conditioning function  $\tau$  with low variance can potentially incur a small excess risk. This makes the minimizer of the variance, a potentially good candidate for conditional meta-learning. We note that  $\text{Var}_\rho(\tau)$  in [\(6\)](#) can be interpreted as a Least-Squares risk associated to the input-(ideal) output pair  $(s, w_\mu)$ . Thanks to this interpretation, we can rely on the following well-known facts, see e.g. [\[11, Lemma A2\]](#).

**Lemma 2** (Best conditioning function in hindsight). *The minimizer of  $\text{Var}_\rho(\cdot)^2$  in [\(6\)](#) over the set  $\mathcal{T}$  is such that  $\tau_\rho(s) = \mathbb{E}_{\mu \sim \rho(\cdot|s)} w_\mu$  almost everywhere on  $\mathcal{S}$ . Moreover, for any  $\tau \in \mathcal{T}$ ,*

$$\text{Var}_\rho(\tau)^2 - \text{Var}_\rho(\tau_\rho)^2 = \mathbb{E}_{s \sim \rho_S} \|\tau(s) - \tau_\rho(s)\|^2. \quad (11)$$

Combining [Thm. 1](#) with [Lemma 2](#), we can formally analyze when the conditional approach is significantly advantageous w.r.t. the unconditional one.

**Conditional vs unconditional meta-learning.** As observed in [\(5\)](#), unconditional meta-learning consists in restricting to the class of constant conditioning functions  $\mathcal{T}^{\text{const}}$ . Minimizing  $\text{Var}_\rho(\cdot)^2$  over this class yields the optimal bias vector for standard meta-learning (see e.g. [\[5, 13, 14, 24\]](#)), given by the expected target tasks' vector  $w_\rho = \mathbb{E}_{\mu \sim \rho_{\mathcal{M}}} w_\mu$ . Applying [\(11\)](#) to the constant function  $\tau \equiv w_\rho$ , we get the following gap between the best performance of conditional and unconditional meta-learning:

$$\text{Var}_\rho(w_\rho)^2 - \text{Var}_\rho(\tau_\rho)^2 = \mathbb{E}_{s \sim \rho_S} \|w_\rho - \tau_\rho(s)\|^2. \quad (12)$$

We note that the gap [\(12\)](#) above is large when the ideal conditioning function  $\tau_\rho$  is ‘‘far’’ from being the constant function  $w_\rho$ . We report below two examples that can be considered illustrative for many real-world scenarios in which such a condition is satisfied. We refer to [App. B](#) for the details and the deduction. In the examples, we parametrize each task with the triplet  $\mu = (w_\mu, \eta_\mu, \xi_\mu)$ , where  $w_\mu$  is the target weight vector,  $\eta_\mu$  is the marginal distribution on the inputs,  $\xi_\mu$  is a noise model and  $y \sim \mu(\cdot|x)$  is  $y = \langle w_\mu, x \rangle + \epsilon$  with  $x \sim \eta_\mu$  and  $\epsilon \sim \xi_\mu$ . Additionally, we denote by  $\mathcal{N}(v, \sigma^2 I)$  a Gaussian distribution with mean  $v \in \mathbb{R}^d$  and covariance matrix  $\sigma^2 I$ , with  $I$  the  $d \times d$  identity matrix.

**Example 1** (Clusters of tasks). *Let  $\rho_{\mathcal{M}} = \frac{1}{m} \sum_{i=1}^m \rho_{\mathcal{M}}^{(i)}$  be a uniform mixture of  $m$  environments (clusters) of tasks. For each  $i = 1, \dots, m$ , a task  $\mu \sim \rho_{\mathcal{M}}^{(i)}$  is sampled such that: 1)  $w_\mu \sim \mathcal{N}(w(i), \sigma_w^2 I)$  with  $w(i) \in \mathbb{R}^d$  a cluster's mean vector and  $\sigma_w^2 I$  a covariance matrix, 2) the marginal  $\eta_\mu = \mathcal{N}(x(i), \sigma_x^2)$  with mean vector  $x(i) \in \mathbb{R}^d$  and variance  $\sigma_x^2$ , 3) the side information is an  $n$  i.i.d. sample from  $\eta_\mu$ , namely  $s = (x_i)_{i=1}^n \sim \eta_\mu^n$ . Then, the gap between conditional and unconditional variance is*

$$\text{Var}_\rho(w_\rho)^2 - \text{Var}_\rho(\tau_\rho)^2 \geq \frac{1}{2m^2} \sum_{i,j=1}^m \left( 1 - \frac{m}{2} e^{-\frac{n}{\sigma_x^2} \|x(i) - x(j)\|^2} \right) \|w(i) - w(j)\|^2. \quad (13)$$

The inequality above confirms our natural intuition. It tells us that the larger is the number of clusters and the more the target weight vectors' and inputs' centroids are distant (i.e. the more the clusters are distant and the inputs' side information are discriminative for conditioning), the more the conditional approach will be advantageous w.r.t. the unconditional one.

**Example 2** (Curve of tasks). Let  $\rho_S$  be a uniform distribution over  $\mathcal{S} = [0, 1]$ . Let  $h : \mathcal{S} \rightarrow \mathbb{R}^d$  parametrize a circle of radius  $r > 0$  centered in  $c \in \mathbb{R}^d$ , such as  $h(s) = r(\cos(2\pi s), \sin(2\pi s), 0, \dots, 0)^\top$ . For  $s \in \mathcal{S}$ , let  $\mu \sim \rho(\cdot|s)$  such that  $w_\mu \sim \mathcal{N}(h(s), \sigma^2 I)$  with  $\sigma \in \mathbb{R}$ . Then,  $\tau_\rho = h$ ,  $w_\rho = c$  and the the gap between conditional and unconditional variance is

$$\text{Var}_\rho(w_\rho)^2 - \text{Var}_\rho(\tau_\rho)^2 = r^2. \quad (14)$$

Hence, in this case, the advantage in applying the conditional approach w.r.t. the unconditional one is equivalent to the squared radius of the circle over which the mean of the target weight vectors  $w_\mu$  lie.

**Conditional meta-learning vs Independent Task Learning (ITL).** Solving each task independently corresponds to choosing the constant conditioning function  $\tau_0 \equiv 0$ . Applying Lemma 2 to this function, the gap between the performance of the best conditional approach and ITL reads as

$$\text{Var}_\rho(0)^2 - \text{Var}_\rho(\tau_\rho)^2 = \mathbb{E}_{s \sim \rho_S} \|w_\rho - \tau_\rho(s)\|^2 + \|w_\rho\|^2. \quad (15)$$

The gap in (15) combines the gain of conditional over unconditional meta-learning with  $\|w_\rho\|^2 = \text{Var}_\rho(0)^2 - \text{Var}_\rho(w_\rho)^2$  that is the advantage of unconditional meta-learning over ITL (see [13, 14]). In the next section, we introduce a convex meta-algorithm mimicking this advantage also in practice.

## 4 Conditional meta-learning algorithm

To address conditional meta-learning in practice, we introduce the following set of conditioning functions. For a given feature map  $\Phi : \mathcal{S} \rightarrow \mathbb{R}^k$  on the side information space, we define the associated space of linear functions

$$\mathcal{T}_\Phi = \left\{ \tau : \mathcal{S} \rightarrow \mathbb{R}^d \mid \tau(\cdot) = M\Phi(\cdot) + b, \text{ for some } M \in \mathbb{R}^{d \times k}, b \in \mathbb{R}^d \right\}. \quad (16)$$

To highlight the dependency of a function  $\tau \in \mathcal{T}_\Phi$  w.r.t. its parameters  $M$  and  $b$ , we will use the notation  $\tau = \tau_{M,b}$ . Evidently,  $\mathcal{T}_\Phi$  contains the space of all unconditional estimators  $\mathcal{T}^{\text{const}}$ . We consider  $\mathcal{T}_\Phi$  equipped with the canonical norm  $\|\tau_{M,b}\|^2 = \|(M, b)\|_F^2 = \|M\|_F^2 + \|b\|^2$ , with  $\|\cdot\|_F$  the Frobenius norm. We now introduce two standard assumptions will allow the design of our method.

**Assumption 2.** The minimizer  $\tau_\rho$  of  $\text{Var}_\rho(\cdot)$  belongs to  $\mathcal{T}_\Phi$ , namely there exist  $M_\rho \in \mathbb{R}^{d \times k}$  and  $b_\rho \in \mathbb{R}^d$ , such that  $\tau_\rho(\cdot) = M_\rho\Phi(\cdot) + b_\rho$ .

**Assumption 3.** There exists  $K > 0$  such that  $\|\Phi(s)\| \leq K$  for any  $s \in \mathcal{S}$ .

Asm. 2 enables us to restrict the conditional meta-learning problem in (5) to  $\mathcal{T}_\Phi$ , rather than to the entire space  $\mathcal{T}$  of measurable functions. In Lemma 7 in App. C we provide the closed forms of  $M_\rho$  and  $b_\rho$  and we express the gap in (12) by the correlation between  $w_\mu$  and  $\Phi(s)$  and the slope of  $\tau_\rho$ . Asm. 3 will allow us to work with a Lipschitz meta-objective, as explained below.

**The convex surrogate problem.** Following a similar strategy to the one adopted for the unconditional setting in [13, 14], we introduce the following surrogate problem for the conditional one in (5):

$$\min_{\tau \in \mathcal{T}} \hat{\mathcal{E}}_\rho(\tau) \quad \hat{\mathcal{E}}_\rho(\tau) = \mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^n} \mathcal{R}_{Z, \tau(s)}(A(\tau(s), Z)), \quad (17)$$

where we have replaced the inner expected risk  $\mathcal{R}_\mu$  with the regularized empirical risk  $\mathcal{R}_{Z, \theta}$  in (2). Exploiting Asm. 2, the problem above can be rewritten more explicitly as follows

$$\min_{M \in \mathbb{R}^{d \times k}, b \in \mathbb{R}^d} \mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^n} \mathcal{L}(M, b, s, Z) \quad \mathcal{L}(M, b, s, Z) = \mathcal{R}_{Z, \tau_{M,b}(s)}(A(\tau_{M,b}(s), Z)). \quad (18)$$

The following proposition characterizes useful properties of the meta-loss  $\mathcal{L}(\cdot, \cdot, s, Z)$  introduced above (such as convexity and differentiability) and it supports its choice as surrogate meta-loss. We denote by  $\cdot^\top$  the standard transposition operation.

**Proposition 3** (Properties of the surrogate meta-loss  $\mathcal{L}$ ). For any  $Z \in \mathcal{D}$  and  $s \in \mathcal{S}$ , the function  $\mathcal{L}(\cdot, \cdot, s, Z)$  is convex, differentiable and its gradient is given by

$$\nabla \mathcal{L}(\cdot, \cdot, s, Z)(M, b) = -\lambda \left( A(\tau_{M,b}(s), Z) - \tau_{M,b}(s) \right) \begin{pmatrix} \Phi(s) \\ 1 \end{pmatrix}^\top \quad (19)$$

for any  $M \in \mathbb{R}^{d \times k}$  and  $b \in \mathbb{R}^d$ . Moreover, under Asm. 1 and Asm. 3, we have

$$\|\nabla \mathcal{L}(\cdot, \cdot, s, Z)(M, b)\|_F^2 \leq L^2 R^2 (K^2 + 1). \quad (20)$$

---

**Algorithm 1** Meta-Algorithm, SGD on (18)

---

**Input**  $\gamma > 0$  meta-step size,  $\lambda > 0$  inner regularization parameter  
**Initialization**  $M_1 = 0 \in \mathbb{R}^{d \times k}$ ,  $b_1 = 0 \in \mathbb{R}^d$   
**For**  $t = 1$  to  $T$   
    Receive  $(\mu_t, s_t) \sim \rho$  and  $Z_t \sim \mu_t^n$   
    Let  $\theta_t = \tau_{M_t, b_t}(s_t) = M_t \Phi(s_t) + b_t$   
    Run the inner algorithm in (2) to obtain  $w_t = A(\theta_t, Z_t)$   
    Compute  $\nabla \mathcal{L}(\cdot, \cdot, s_t, Z_t)(M_t, b_t) = -\lambda(w_t - \theta_t) \begin{pmatrix} \Phi(s_t) \\ \mathbf{1} \end{pmatrix}^\top$  as in (19)  
    Update  $(M_{t+1}, b_{t+1}) = (M_t, b_t) - \gamma \nabla \mathcal{L}(\cdot, \cdot, s_t, Z_t)(M_t, b_t)$   
**Return**  $\bar{M} = \frac{1}{T} \sum_{t=1}^T M_t$ ,  $\bar{b} = \frac{1}{T} \sum_{t=1}^T b_t$

---

The proof of Prop. 3 is reported in App. D.1 and it follows a similar reasoning in [14], by taking into account also the parameter  $M$  in the optimization problem.

**The conditional meta-learning estimator.** In this work we propose to apply Stochastic Gradient Descent (SGD) on the surrogate problem in (18). Alg. 1 summarizes the implementation of this approach: assuming a sequence of i.i.d. pairs  $(Z_t, s_t)_{t=1}^T$  of training sets and side information, at each iteration the algorithm updates the conditional iterates  $(M_t, b_t)$  by performing a step of constant size  $\gamma > 0$  in the direction of  $-\nabla \mathcal{L}(\cdot, \cdot, s_t, Z_t)(M_t, b_t)$ . The map  $\tau_{\bar{M}, \bar{b}}$  is then returned as conditional estimator, with  $(\bar{M}, \bar{b})$  the average across all the iterates  $(M_t, b_t)_{t=1}^T$ . The following result characterizes the excess risk of the proposed estimator.

**Theorem 4** (Excess risk bound for the conditioning function returned by Alg. 1). *Let Asm. 1 and Asm. 3 hold. Let  $\tau_{M, b}$  be a fixed function in  $\mathcal{T}_\Phi$  and let  $\text{Var}_\rho(\tau_{M, b})^2$  be the corresponding variance introduced in (7). Let  $\bar{M}$  and  $\bar{b}$  be the outputs of Alg. 1 applied to a sequence  $(Z_t, s_t)_{t=1}^T$  of i.i.d. pairs sampled from  $\rho$  with inner regularization parameter and meta-step size*

$$\lambda = \frac{2RL}{\text{Var}_\rho(\tau_{M, b})} \frac{1}{\sqrt{n}} \quad \gamma = \frac{\|(M, b)\|_F}{LR\sqrt{K^2 + 1}} \frac{1}{\sqrt{T}}. \quad (21)$$

Then, in expectation w.r.t. the sampling of  $(Z_t, s_t)_{t=1}^T$ ,

$$\mathbb{E} \mathcal{E}_\rho(\tau_{\bar{M}, \bar{b}}) - \mathcal{E}_\rho^* \leq \frac{2RL\text{Var}_\rho(\tau_{M, b})}{\sqrt{n}} + \frac{LR\sqrt{K^2 + 1} \|(M, b)\|_F}{\sqrt{T}}. \quad (22)$$

**Proof (Sketch).** We consider the following decomposition

$$\mathbb{E} \mathcal{E}_\rho(\tau_{\bar{M}, \bar{b}}) - \mathcal{E}_\rho^* = \underbrace{\mathbb{E} \mathcal{E}_\rho(\tau_{\bar{M}, \bar{b}}) - \hat{\mathcal{E}}_\rho(\tau_{\bar{M}, \bar{b}})}_{\text{B}} + \underbrace{\mathbb{E} \hat{\mathcal{E}}_\rho(\tau_{\bar{M}, \bar{b}}) - \hat{\mathcal{E}}_\rho(\tau_{M, b})}_{\text{C}} + \underbrace{\hat{\mathcal{E}}_\rho(\tau_{M, b}) - \mathcal{E}_\rho^*}_{\text{D}}. \quad (23)$$

Applying Asm. 1 and the stability arguments in Prop. 5 in App. A, we can write  $\text{B} \leq 2R^2L^2(\lambda n)^{-1}$ . The term C is the term expressing the convergence rate of Alg. 1 on the surrogate problem in (18) and, exploiting Asm. 3 and Prop. 3, it can be controlled as described in Prop. 9 in App. D.2. Regarding the term D, exploiting the definition of the algorithm in (2), we can write  $\text{D} \leq \frac{\lambda}{2} \text{Var}_\rho(\tau_{M, b})^2$ . Combining all the terms and optimizing w.r.t.  $\gamma$  and  $\lambda$ , we get the desired statement. ■

We now comment about the result we got above in Thm. 4.

**Proposed vs optimal conditioning function.** Specializing the bound in Thm. 4 to the best conditioning function  $\tau_\rho$  in Lemma 2, thanks to Asm. 2, we get the following bound for our estimator:

$$\mathbb{E} \mathcal{E}_\rho(\tau_{\bar{M}, \bar{b}}) - \mathcal{E}_\rho^* \leq \mathcal{O}\left(\text{Var}_\rho(\tau_\rho) n^{-1/2} + \|(M_\rho, b_\rho)\|_F T^{-1/2}\right). \quad (24)$$

Hence, our proposed meta-algorithm achieves comparable performance to the best conditioning function  $\tau_\rho$  in hindsight, provided that the number of observed tasks is sufficiently large. The bound

above also highlights the trade-off between statistical and computational complexity of the class  $\mathcal{T}_\Phi$ : conditional meta-learning incurs in a cost  $\|(M_\rho, b_\rho)\|_F$  in the  $\sqrt{T}$ -term that is larger than the  $\|b_\rho\|$  cost of unconditional meta-learning (see [5, 13, 24]), which is, however, limited to constant conditioning functions. This is an acceptable price, since, as we discussed in Sec. 3, the performance of conditional meta-learning is significantly better than the standard one in many common scenarios.

**Remark 3.** When  $\tau_\rho \notin \mathcal{T}_\Phi$  (i.e. when *Asm. 3* does not hold), our method suffers an additional approximation error due to the fact  $\min_{\tau \in \mathcal{T}_\Phi} \text{Var}_\rho(\tau) > \text{Var}_\rho(\tau_\rho)$ . In this case, one might nullify the gap above by considering a feature map  $\Phi : \mathcal{S} \rightarrow \mathcal{H}$  with  $\mathcal{H}$  a universal reproducing kernel Hilbert space of functions. Exploiting standard arguments from online learning with kernels literature (see e.g. [25, 36, 37]), in *Lemma 10* in *App. D.3* we describe the implementation of *Alg. 1* for this setting using only evaluations of the kernel associated to the feature map. We leave the corresponding theoretical analysis to future work.

**Proposed conditioning function vs unconditional meta-learning.** Specializing *Thm. 4* to  $\tau_{M,b} \equiv w_\rho$ , the bound for our estimator becomes:

$$\mathbb{E} \mathcal{E}_\rho(\tau_{\bar{M}, \bar{b}}) - \mathcal{E}_\rho^* \leq \mathcal{O}(\text{Var}_\rho(w_\rho) n^{-1/2} + \|w_\rho\| T^{-1/2}), \quad (25)$$

which is equivalent to state-of-the-art bounds for unconditional methods, see [5, 13, 14, 24]. Hence, our conditional approach provides, at least, the same guarantees as its unconditional counterpart.

**Proposed conditioning function vs ITL.** Specializing *Thm. 4* to  $\tau_{M,b} \equiv 0$  corresponds to force  $\gamma = 0$  and, consequently, *Alg. 1* to not move. In such a case, we get the bound:

$$\mathbb{E} \mathcal{E}_\rho(\tau_{\bar{M}, \bar{b}}) - \mathcal{E}_\rho^* \leq \mathcal{O}(\text{Var}_\rho(0) n^{-1/2}), \quad (26)$$

which corresponds to the standard excess risk bound for ITL, see [5, 13, 14, 24]. In other words, our method does not generate negative transfer effect.

**Remark 4** (Fine tuning). In the case of the online inner family in *Rem. 1* used in fine tuning, *Alg. 1* employs an approximation of the meta-subgradient in (19) by replacing the batch regularized empirical risk minimizer  $A(\tau_{M,b}(s), Z)$  in (2) with the last iterate of the online algorithm in (3). As shown in [13, 14] for the unconditional setting, such an approximation does not affect the behavior of the bounds above.

## 5 Experiments

In this section we compare the numerical performance<sup>1</sup> of our conditional method in *Alg. 1* (cond.) w.r.t. its unconditional counterpart in [13] (uncond.). We will also add to the comparison the methods consisting in applying the inner algorithm on each task with  $\tau \equiv 0 \in \mathbb{R}^d$  (i.e. ITL) and the unconditional oracle  $\tau \equiv w_\rho = \mathbb{E}_{\mu \sim \rho_{\mathcal{M}}} w_\mu$  (mean), when available. We considered regression problems and we evaluated the errors by the absolute loss. The results refer to the fine tuning variant of the methods with the online inner algorithm in (3). In *App. E*, we describe how we tuned the hyper-parameters  $\lambda, \gamma$  in our experiments.

**Synthetic clusters.** We considered three variants of the setting described in *Ex. 1*. In all the variants we sampled  $T_{\text{tot}} = 480$  tasks from a mixture of  $m$  clusters with the same probability. For each task  $\mu$ , we sampled the corresponding target vector  $w_\mu$  from the  $d = 20$ -dimensional Gaussian distribution  $\mathcal{N}(w(j_\mu), I)$ , where,  $j_\mu \in \{1, \dots, m\}$  denotes the cluster from which the task  $\mu$  was sampled. We then generated the corresponding dataset  $(x_i, y_i)_{i=1}^{n_{\text{tot}}}$  with  $n_{\text{tot}} = 20$ . We sampled the inputs from  $\mathcal{N}(x(j_\mu), I)$  and we generated the labels according to the equation  $y = \langle x, w_\mu \rangle + \epsilon$ , with the noise  $\epsilon$  sampled from  $\mathcal{N}(0, \sigma^2 I)$ , with  $\sigma$  chosen in order to have signal-to-noise ratio equal to 1.

We implemented our conditional method using as side information the training input points  $X = (x_i)_{i=1}^{n_{\text{tr}}} \in \bigcup_{n \in \mathbb{N}} \mathcal{X}^n$  and the feature map  $\Phi : \bigcup_{n \in \mathbb{N}} \mathcal{X}^n \rightarrow \mathbb{R}^d$  defined by  $\Phi(X) = \frac{1}{n_{\text{tr}}} \sum_{i=1}^{n_{\text{tr}}} x_i$ .

In *Fig. 1* (left-top), we generated an environment as above with just one cluster ( $m = 1$ ) and we took  $w(1) = 4 \in \mathbb{R}^d$  (the vector in  $\mathbb{R}^d$  with all components 4) and  $x(1) = 1 \in \mathbb{R}^d$ . As we can see, coherently with previous work [13], the unconditional approach outperforms ITL and it converges to the mean vector  $w_\rho = w(1)$  as the number of training tasks increases. The conditional approach returns equivalent performances to the unconditional counterpart.

<sup>1</sup>Code is available at <https://github.com/dGiulia/ConditionalMetaLearning.git>

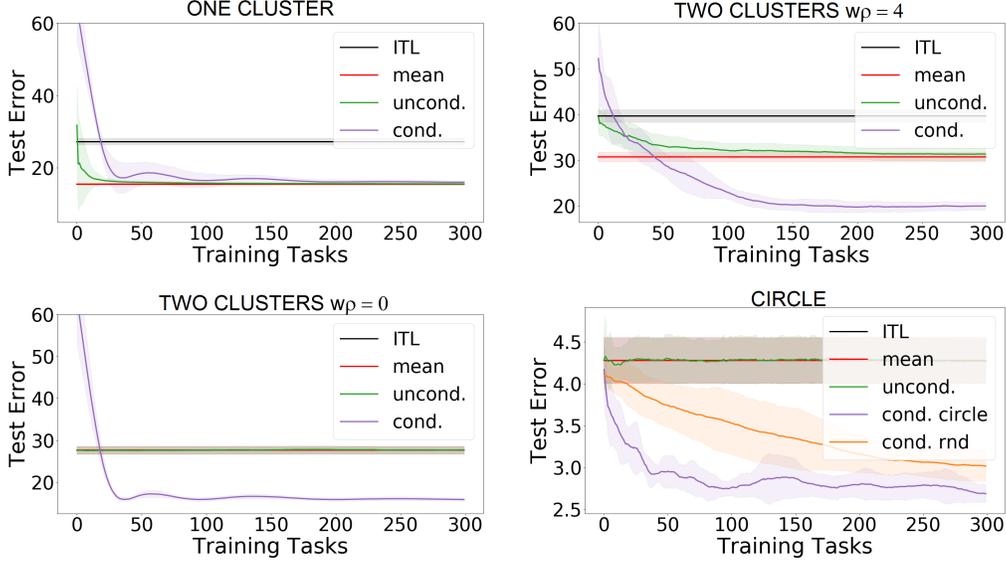


Figure 1: Performance (averaged over 10 seeds) of different methods w.r.t. an increasing number of tasks on different environments: with one cluster (left-top), with two clusters and  $w_\rho = 4$  (right-top), with two clusters and  $w_\rho = 0$  (left-bottom), circle (right-bottom).

In Fig. 1 (right-top), we considered an environment of two clusters ( $m = 2$ ) identified by  $w(1) = 8 \in \mathbb{R}^d$ ,  $w(2) = 0 \in \mathbb{R}^d$  (implying  $w_\rho = 4$ ),  $x(1) = 1 \in \mathbb{R}^d$  and  $x(2) = -x(1)$ . As we can see, the conditional approach outperform ITL as in the previous setting, but the conditional approach yields even better performance.

Finally, in Fig. 1 (left-bottom), we considered an environment of two clusters ( $m = 2$ ) identified by  $w(1) = 4 \in \mathbb{R}^d$ ,  $w(2) = -w(1)$  (implying  $w_\rho = 0$ ),  $x(1) = 1 \in \mathbb{R}^d$  and  $x(2) = -x(1)$ . As expected, the unconditional approach mimics the poor performance of ITL, while, the performance of the conditional approach is promising.

Summarizing, the conditional approach brings advantage w.r.t. the unconditional one when the heterogeneity of the environment is significant. When the environment is homogeneous, the performance of the two are equivalent. This conclusion is exactly inline with our theory in (25) and (26).

**Synthetic circle.** We sampled  $T_{\text{tot}} = 480$  tasks according to the setting described in Ex. 2. Specifically, for each task  $\mu$ , we first sampled the corresponding side information  $s \in [0, 1]$  according to the uniform distribution. We then generated the vector  $h(s) = r (\cos(2\pi s), \sin(2\pi s), 0, \dots, 0)^\top \in \mathbb{R}^d$ , with  $d = 20$ , on the zero-centered circle of radius  $r = 8$ . After this, we sampled the corresponding target weight vector  $w_\mu$  from  $\mathcal{N}(h(s), I)$ . We then generated the associated dataset of  $n_{\text{tot}} = 20$  points as for the experiments above. We applied our conditional approach with the true underlying feature map  $\Phi(s) = (\cos(2\pi s), \sin(2\pi s))$  (cond. circle) and the feature map mimicking a Gaussian distribution by Fourier random features [34] described below (at the end of this section) with parameters  $k = 50$  and  $\sigma = 10$  (cond. rnd).

From Fig. 1 (right-bottom) we see that the performance of unconditional meta-learning mimics the poor performance of ITL (in fact, we have  $w_\rho = 0$ ). On the other hand, both the conditional approaches bring a substantial advantage and the random features' variant approaches the variant knowing the true underlying feature map.

**Lenk dataset.** We considered the computer survey data from [26, 30], in which  $T_{\text{tot}} = 180$  people (tasks) rated the likelihood of purchasing one of  $n_{\text{tot}} = 20$  different personal computers. The input represents  $d = 13$  different computers' characteristics, while the output is an integer rating from 0 to 10. In this case, we used as side information the training datapoints  $Z = (z_i)_{i=1}^{n_{\text{tr}}}$  and the feature map  $\Phi : \mathcal{D} \rightarrow \mathbb{R}^{2d}$  defined by  $\Phi(Z) = \frac{1}{n_{\text{tr}}} \sum_{i=1}^{n_{\text{tr}}} \phi(z_i)$ , with  $\phi(z_i) = \text{vec}(x_i(y_i, 1)^\top)$ , where, for any matrix  $A = [a_1, a_2] \in \mathbb{R}^{d \times 2}$  with columns  $a_1, a_2 \in \mathbb{R}^d$ ,  $\text{vec}(A) = (a_1, a_2)^\top \in \mathbb{R}^{2d}$ . Fig. 2 (left) shows that, coherently to previous literature [13], the unconditional approach significantly outperforms ITL, but the performance of its conditional counterpart is even better.

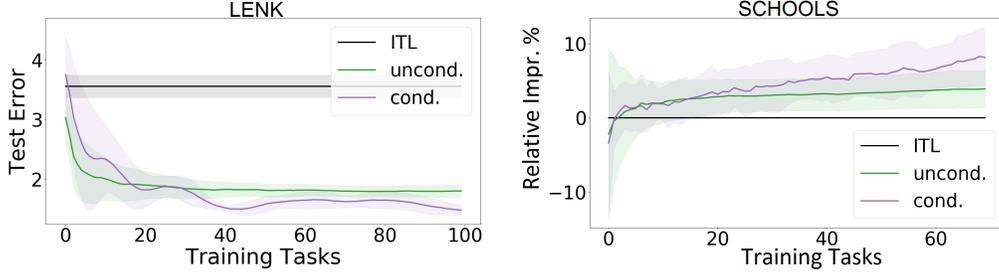


Figure 2: Performance (averaged over 10 seeds) of different methods w.r.t. an increasing number of tasks. Lenk dataset (left), Schools dataset (right).

**Schools dataset.** We considered the Schools dataset [3], consisting of examination records from  $T_{\text{tot}} = 139$  schools. Each school is associated to a task, individual students are represented by a features’ vectors  $x \in \mathbb{R}^d$ , with  $d = 26$ , and their exam scores to the outputs. The sample size  $n_{\text{tot}}$  varies across the tasks from a minimum 24 to a maximum 251. We used as side information the training inputs  $X = (x_i)_{i=1}^{n_{\text{tr}}}$  and the feature map mimicking a Gaussian distribution by Fourier random features described below (at the end of this section) with parameters  $k = 1000$  and  $\sigma = 100$ . Fig. 2 (right) shows that, also in this case, the unconditional approach brings a meaningful improvement w.r.t. ITL, but the gain provided by its conditional counterpart is even more evident.

**Feature map by Fourier random features.** We now describe the feature map mimicking a Gaussian distribution by Fourier random features [34] we used in our synthetic circle experiment and Schools dataset experiment. We recall that, in these cases, we considered as side information the inputs  $X = (x_i)_{i=1}^n$ . The feature map above was then defined as  $\Phi(X) = \frac{1}{n} \sum_{i=1}^n \phi(x_i)$ , where,  $\phi$  was built as follows. We first introduced an integer  $k \in \mathbb{N}$  and a constant  $\sigma \in \mathbb{R}$ . We then sampled a vector  $v \in \mathbb{R}^k$  from the uniform distribution over  $[0, 2\pi]^k$  and a matrix  $U \in \mathbb{R}^{k \times d}$  is sampled from the Gaussian distribution  $\mathcal{N}(0, \sigma I)$ . We then defined

$$\phi(x_i) = \sqrt{\frac{2}{k}} \cos(Ux_i + v) \in \mathbb{R}^k, \quad (27)$$

where  $\cos(\cdot)$  is applied component-wise to the vector.

## 6 Conclusion

We proposed a new conditional meta-learning framework for biased regularization and fine tuning based on side information and we provided a theoretical analysis demonstrating its potential advantage over standard meta-learning, when the environment of tasks is heterogeneous. In the future, taking inspiration from [12, 32], it would be interesting to develop a variant of our method in which the hyper-parameters are automatically tuned in efficient way. In addition, it would valuable to extend our conditional approach and the corresponding analysis to other meta-learning paradigms considering different families of inner algorithms, such as [14, 38].

## Broader impact

Meta-learning is a very important field for machine learning with potential societal implications related to applications such as recommender systems. In this work we focused mostly on theoretical and modeling aspects, however in the future the topic will need to take into consideration contributions from other fields related to ethical and societal aspects, such as privacy and fairness.

## Acknowledgments and Disclosure of Funding

This work was supported in part by SAP SE and EPSRC Grant N. EP/P009069/1. C.C. acknowledges the Royal Society (grant SPREM RGS\R1\201149).

## References

- [1] J. Abernethy, F. Bach, T. Evgeniou, and J.-P. Vert. A new approach to collaborative filtering: Operator estimation with spectral regularization. *Journal of Machine Learning Research*, 10(Mar):803–826, 2009.
- [2] A. Argyriou, S. Cléménçon, and R. Zhang. Learning the graph of relations among multiple tasks. 2013.
- [3] A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.
- [4] A. Argyriou, A. Maurer, and M. Pontil. An algorithm for transfer learning in a heterogeneous environment. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 71–85. Springer, 2008.
- [5] M.-F. Balcan, M. Khodak, and A. Talwalkar. Provable guarantees for gradient-based meta-learning. In *International Conference on Machine Learning*, pages 424–433, 2019.
- [6] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator theory in Hilbert Spaces*, volume 408. Springer, 2011.
- [7] J. Baxter. A model of inductive bias learning. *J. Artif. Intell. Res.*, 12(149–198):3, 2000.
- [8] O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.
- [9] T. T. Cai, T. Liang, and A. Rakhlin. Weighted message passing and minimum energy flow for heterogeneous stochastic block models with side information. *Journal of Machine Learning Research*, 21(11):1–34, 2020.
- [10] G. Cavallanti, N. Cesa-Bianchi, and C. Gentile. Linear algorithms for online multitask classification. *Journal of Machine Learning Research*, 11:2901–2934, 2010.
- [11] C. Ciliberto, L. Rosasco, and A. Rudi. A general framework for consistent structured prediction with implicit loss embeddings. *arXiv preprint arXiv:2002.05424*, 2020.
- [12] A. Cutkosky and F. Orabona. Black-box reductions for parameter-free online learning in banach spaces. In *Proceedings of the 31st Conference On Learning Theory*, volume 75, 2018.
- [13] G. Denevi, C. Ciliberto, R. Grazi, and M. Pontil. Learning-to-learn stochastic gradient descent with biased regularization. In *International Conference on Machine Learning*, pages 1566–1575, 2019.
- [14] G. Denevi, D. Stamos, C. Ciliberto, and M. Pontil. Online-within-online meta-learning. In *Advances in Neural Information Processing Systems*, pages 13089–13099, 2019.
- [15] T. Evgeniou and M. Pontil. Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 109–117, 2004.
- [16] A. Fallah, A. Mokhtari, and A. Ozdaglar. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. *arXiv preprint arXiv:1908.10400*, 2019.
- [17] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR, 2017.
- [18] C. Finn and S. Levine. Meta-learning and universality: Deep representations and gradient descent can approximate any learning algorithm. In *International Conference on Learning Representations*, 2018.
- [19] C. Finn, A. Rajeswaran, S. Kakade, and S. Levine. Online meta-learning. In *International Conference on Machine Learning*, pages 1920–1930, 2019.
- [20] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- [21] L. Jacob, J.-p. Vert, and F. R. Bach. Clustered multi-task learning: A convex formulation. In *Advances in neural information processing systems*, pages 745–752, 2009.
- [22] G. Jerfel, E. Grant, T. Griffiths, and K. A. Heller. Reconciling meta-learning and continual learning with online mixtures of tasks. In *Advances in Neural Information Processing Systems*, pages 9119–9130, 2019.

- [23] K. Ji, J. Yang, and Y. Liang. Multi-step model-agnostic meta-learning: Convergence and improved algorithms. *arXiv preprint arXiv:2002.07836*, 2020.
- [24] M. Khodak, M.-F. F. Balcan, and A. S. Talwalkar. Adaptive gradient-based meta-learning methods. In *Advances in Neural Information Processing Systems*, pages 5915–5926, 2019.
- [25] J. Kivinen, A. J. Smola, and R. C. Williamson. Online learning with kernels. *IEEE transactions on signal processing*, 52(8):2165–2176, 2004.
- [26] P. J. Lenk, W. S. DeSarbo, P. E. Green, and M. R. Young. Hierarchical bayes conjoint analysis: Recovery of partworth heterogeneity from reduced experimental designs. *Marketing Science*, 15(2):173–191, 1996.
- [27] A. Maurer. The rademacher complexity of linear transformation classes. In *International Conference on Computational Learning Theory*, pages 65–78, 2006.
- [28] A. Maurer. Transfer bounds for linear feature learning. *Machine Learning*, 75(3):327–350, 2009.
- [29] A. Maurer and M. Pontil. Transfer learning in a heterogeneous environment. In *2012 3rd International Workshop on Cognitive Information Processing (CIP)*, pages 1–6. IEEE, 2012.
- [30] A. M. McDonald, M. Pontil, and D. Stamos. New perspectives on k-support and cluster norms. *Journal of Machine Learning Research*, 17(155):1–38, 2016.
- [31] A. Nichol, J. Achiam, and J. Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [32] F. Orabona and D. Pál. Coin betting and parameter-free online learning. In *Advances in Neural Information Processing Systems*, pages 577–585, 2016.
- [33] A. Pentina and C. Lampert. A PAC-Bayesian bound for lifelong learning. In *International Conference on Machine Learning*, pages 991–999, 2014.
- [34] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2008.
- [35] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018.
- [36] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- [37] A. Singh, N. Ahuja, and P. Moulin. Online learning with kernels: Overcoming the growing sum problem. In *2012 IEEE International Workshop on Machine Learning for Signal Processing*, pages 1–6. IEEE, 2012.
- [38] N. Tripuraneni, C. Jin, and M. I. Jordan. Provable meta-learning of linear representations. *arXiv preprint arXiv:2002.11684*, 2020.
- [39] R. Vuorio, S.-H. Sun, H. Hu, and J. J. Lim. Multimodal model-agnostic meta-learning via task-aware modulation. In *Advances in Neural Information Processing Systems*, pages 1–12, 2019.
- [40] R. Wang, Y. Demiris, and C. Ciliberto. A structured prediction approach for conditional meta-learning. *Advances in Neural Information Processing Systems*, 2020.
- [41] H. Yao, Y. Wei, J. Huang, and Z. Li. Hierarchically structured meta-learning. *arXiv preprint arXiv:1905.05301*, 2019.

## Appendix

The supplementary material is organized as follows. In [App. A](#) we give the bound on the generalization error of the algorithm in (2) that we used in various proofs. In [App. B](#) we formally describe the deduction of the statements reported in [Ex. 1](#) and [Ex. 2](#) in [Sec. 3](#). In [App. C](#) we report the closed form of  $M_\rho \in \mathbb{R}^{d \times k}$  and  $b_\rho \in \mathbb{R}^d$  in [Asm. 2](#) and we express the gap between the conditional and the unconditional variance in (12) by the correlation between the target tasks' vectors  $w_\mu$  and the transformed side information  $\Phi(s)$  or the slope of  $\tau_\rho$ . In [App. D](#), we report the proofs of the statements we used in [Sec. 4](#) in order to prove the expected excess risk bound in [Thm. 4](#) for [Alg. 1](#). Finally, in [App. E](#), we report the experimental details we omitted in the main body.

### A Generalization error of the algorithm in (2)

In this section we report the generalization error bound of the family of inner algorithms in (2) that we used in our proofs. The statement exploits standard tools from stability theory. We do not claim any originality, we report the proof for completeness.

**Proposition 5** (Generalization error of the algorithm in (2)). *For a distribution  $\mu \sim \rho$ , fix a dataset  $Z = (x_i, y_i)_{i=1}^n \sim \mu^n$  and, for any  $i \in \{1, \dots, n\}$ , fix a datapoint  $z'_i = (x'_i, y'_i) \sim \mu$  independent from  $Z$ . For any  $\theta \in \Theta$  not depending on  $Z$ , let  $\hat{w}_\theta(Z) = A(\theta, Z)$  be the output of the algorithm in (2) over  $Z$  and let  $s'_{\theta,i} \in \partial \ell(\cdot, y'_i)(\langle x'_i, \hat{w}_\theta(Z) \rangle)$  be a subgradient of  $\ell(\cdot, y'_i)$  at  $\langle x'_i, \hat{w}_\theta(Z) \rangle$ . Then, the following generalization error bound holds for  $\hat{w}_\theta(Z)$*

$$\mathbb{E}_{Z \sim \mu^n} [\mathcal{R}_\mu(\hat{w}_\theta(Z)) - \mathcal{R}_Z(\hat{w}_\theta(Z))] \leq \frac{2}{\lambda n} \mathbb{E}_{Z \sim \mu^n} \mathbb{E}_{z'_i \sim \mu} \|x'_i s'_{\theta,i}\|^2. \quad (28)$$

As a consequence, under [Asm. 1](#), the right side term above can be upper bounded by  $2L^2 R^2 (\lambda n)^{-1}$ .

**Proof.** For any  $i \in \{1, \dots, n\}$ , consider the dataset  $Z^{(i)}$ , a copy of the original dataset  $Z$  in which we exchange the point  $z_i = (x_i, y_i)$  with the new i.i.d. point  $z'_i = (x'_i, y'_i)$ . For a fixed  $\theta \in \Theta$ , we analyze how much this perturbation affects the outputs of the algorithm in (2). In other words, we study the discrepancy between  $\hat{w}_\theta(Z)$  and  $\hat{w}_\theta(Z^{(i)})$ . We start from observing that, since  $\mathcal{R}_{Z,\theta}$  is  $\lambda$ -strongly convex w.r.t.  $\|\cdot\|$ , by growth condition and the definition of the algorithm in (2), we can write the following

$$\begin{aligned} \frac{\lambda}{2} \|\hat{w}_\theta(Z^{(i)}) - \hat{w}_\theta(Z)\|^2 &\leq \mathcal{R}_{Z,\theta}(\hat{w}_\theta(Z^{(i)})) - \mathcal{R}_{Z,\theta}(\hat{w}_\theta(Z)) \\ \frac{\lambda}{2} \|\hat{w}_\theta(Z^{(i)}) - \hat{w}_\theta(Z)\|^2 &\leq \mathcal{R}_{Z^{(i)},\theta}(\hat{w}_\theta(Z)) - \mathcal{R}_{Z^{(i)},\theta}(\hat{w}_\theta(Z^{(i)})). \end{aligned} \quad (29)$$

Hence, summing the two inequalities above, we get

$$\begin{aligned} \lambda \|\hat{w}_\theta(Z^{(i)}) - \hat{w}_\theta(Z)\|^2 &\leq \mathcal{R}_{Z,\theta}(\hat{w}_\theta(Z^{(i)})) - \mathcal{R}_{Z^{(i)},\theta}(\hat{w}_\theta(Z^{(i)})) + \mathcal{R}_{Z^{(i)},\theta}(\hat{w}_\theta(Z)) - \mathcal{R}_{Z,\theta}(\hat{w}_\theta(Z)) \\ &= \frac{\mathbf{B} + \mathbf{C}}{n}, \end{aligned} \quad (30)$$

where we have introduced the terms

$$\begin{aligned} \mathbf{B} &= \ell(\langle x'_i, \hat{w}_\theta(Z) \rangle, y'_i) - \ell(\langle x'_i, \hat{w}_\theta(Z^{(i)}) \rangle, y'_i) \\ \mathbf{C} &= \ell(\langle x_i, \hat{w}_\theta(Z^{(i)}) \rangle, y_i) - \ell(\langle x_i, \hat{w}_\theta(Z) \rangle, y_i). \end{aligned} \quad (31)$$

Now, exploiting the assumption  $s'_{\theta,i} \in \partial \ell(\cdot, y'_i)(\langle x'_i, \hat{w}_\theta(Z) \rangle)$ , applying Holder's inequality and introducing a subgradient  $s_{\theta,i} \in \partial \ell(\cdot, y_i)(\langle x_i, \hat{w}_\theta(Z^{(i)}) \rangle)$ , we can write

$$\begin{aligned} \mathbf{B} &\leq \langle x'_i s'_{\theta,i}, \hat{w}_\theta(Z) - \hat{w}_\theta(Z^{(i)}) \rangle \leq \|x'_i s'_{\theta,i}\| \|\hat{w}_\theta(Z^{(i)}) - \hat{w}_\theta(Z)\| \\ \mathbf{C} &\leq \langle x_i s_{\theta,i}, \hat{w}_\theta(Z^{(i)}) - \hat{w}_\theta(Z) \rangle \leq \|x_i s_{\theta,i}\| \|\hat{w}_\theta(Z^{(i)}) - \hat{w}_\theta(Z)\|. \end{aligned} \quad (32)$$

Combining these last two inequalities with (30) and simplifying, we get the following

$$\|\hat{w}_\theta(Z^{(i)}) - \hat{w}_\theta(Z)\| \leq \frac{1}{\lambda n} \left( \|x'_i s'_{\theta,i}\| + \|x_i s_{\theta,i}\| \right). \quad (33)$$

Hence, combining the first row in (32) with (33), we can write

$$\ell(\langle x'_i, \hat{w}_\theta(Z) \rangle, y'_i) - \ell(\langle x'_i, \hat{w}_\theta(Z^{(i)}) \rangle, y'_i) \leq \frac{1}{\lambda n} \left( \|x'_i s'_{\theta,i}\|^2 + \|x'_i s'_{\theta,i}\| \|x_i s_{\theta,i}\| \right). \quad (34)$$

Now, taking the expectation w.r.t.  $Z \sim \mu^n$  and  $z'_i \sim \mu$  of the left side member above, according to [8, Lemma 7], we get

$$\mathbb{E}_{Z \sim \mu^n} \mathbb{E}_{z'_i \sim \mu} \left[ \ell(\langle x'_i, \hat{w}_\theta(Z) \rangle, y'_i) - \ell(\langle x'_i, \hat{w}_\theta(Z^{(i)}) \rangle, y'_i) \right] = \mathbb{E}_{Z \sim \mu^n} \left[ \mathcal{R}_\mu(\hat{w}_\theta(Z)) - \mathcal{R}_Z(\hat{w}_\theta(Z)) \right].$$

Finally, taking the expectation of the right side member, exploiting the fact that the points are i.i.d. according  $\mu$ , we get

$$\mathbb{E}_{Z \sim \mu^n} \mathbb{E}_{z'_i \sim \mu} \frac{1}{\lambda n} \left( \|x'_i s'_{\theta,i}\|^2 + \|x'_i s'_{\theta,i}\| \|x_i s_{\theta,i}\| \right) \leq \frac{2}{\lambda n} \mathbb{E}_{Z \sim \mu^n} \mathbb{E}_{z'_i \sim \mu} \|x'_i s'_{\theta,i}\|^2, \quad (35)$$

where we recall that  $s'_{\theta,i} \in \partial \ell(\cdot, y'_i)(\langle x'_i, \hat{w}_\theta(Z) \rangle)$ . The statement derives from combining the two last statements above with the expectation w.r.t.  $Z \sim \mu^n$  and  $z'_i \sim \mu$  of (34). The second statement directly derives from the first one, once one observes that, if  $\ell(\cdot, y)$  is  $L$ -Lipschitz for any  $y \in \mathcal{Y}$ , then,  $|s'_{\theta,i}| \leq L$  (see [36, Lemma 14.7]).  $\blacksquare$

## B Examples

In this section, we provide the deduction of the statements in the examples reported in Sec. 3. We start from presenting some computation regarding a generic environment parametrized by a latent variable in App. B.1 and, then, in App. B.2, we specify this computation and we derive the statement in Ex. 1. Finally, in App. B.3, we prove the statement in Ex. 2.

### B.1 General parametrization

Consider the case where a latent variable  $\alpha \in \mathcal{A}$  parametrizes the environment  $\rho$ . Denote by  $\rho(\cdot|\alpha)$  the conditional distributions given  $\alpha$  and by  $\rho_{\mathcal{A}}$  the marginal distribution of the latent variable. As usual, we assume  $\rho(\mu, \alpha) = \rho(\mu|\alpha)\rho_{\mathcal{A}}(\alpha)$ . Introduce also

$$w(\alpha) = \int w_\mu \rho(\mu|\alpha) dw_\mu \quad \sigma(\alpha)^2 = \frac{1}{2} \int \|w_\mu - w(\alpha)\|^2 \rho(\mu|\alpha) dw_\mu \quad (36)$$

the conditional expectation and the conditional variance of the target weight vectors  $w_\mu$  given  $\alpha$ , respectively. We now explicitly compute the unconditional and the conditional variance for this generic environment.

**Unconditional variance.** We start from observing that thanks to the parametrization of the environment  $\rho$ , we can rewrite the unconditional variance as follows

$$\begin{aligned} \text{Var}_\rho(w_\rho)^2 &= \mathbb{E}_{\mu \sim \rho_{\mathcal{M}}} \|w_\mu - w_\rho\|^2 = \int \|w_\mu - w_\rho\|^2 \rho(\mu) dw_\mu \\ &= \int \left( \int \|w_\mu - w_\rho\|^2 \rho(\mu|\alpha) dw_\mu \right) \rho_{\mathcal{A}}(\alpha) d\alpha. \end{aligned} \quad (37)$$

We now observe that, for any  $\alpha \in \mathcal{A}$ , we can write the following

$$\begin{aligned} &\int \|w_\mu - w_\rho\|^2 \rho(\mu|\alpha) dw_\mu \\ &= \int \left( \|w_\mu\|^2 - 2 \langle w_\mu, w_\rho \rangle + \|w_\rho\|^2 \right) \rho(\mu|\alpha) dw_\mu \\ &= \int \|w_\mu\|^2 \rho(\mu|\alpha) dw_\mu - 2 \langle w(\alpha), w_\rho \rangle + \|w_\rho\|^2 \\ &= \int \|w_\mu\|^2 \rho(\mu|\alpha) dw_\mu \pm \|w(\alpha)\|^2 - 2 \langle w(\alpha), w_\rho \rangle + \|w_\rho\|^2 \\ &= \int \|w_\mu - w(\alpha)\|^2 \rho(\mu|\alpha) dw_\mu + \|w(\alpha) - w_\rho\|^2 \\ &= 2\sigma(\alpha)^2 + \|w(\alpha) - w_\rho\|^2. \end{aligned} \quad (38)$$

Hence, substituting in (37), we get

$$\text{Var}_\rho(w_\rho)^2 = 2 \int \sigma(\alpha)^2 \rho_{\mathcal{A}}(\alpha) d\alpha + \int \|w(\alpha) - w_\rho\|^2 \rho_{\mathcal{A}}(\alpha) d\alpha. \quad (39)$$

We now observe that the second term above can be rewritten as follows

$$\begin{aligned} & \int \|w(\alpha) - w_\rho\|^2 \rho_{\mathcal{A}}(\alpha) d\alpha \\ &= \int \|w(\alpha)\|^2 \rho_{\mathcal{A}}(\alpha) d\alpha - \|w_\rho\|^2 \\ &= \int \|w(\alpha)\|^2 \rho_{\mathcal{A}}(\alpha) d\alpha - \left\| \int w(\alpha') \rho_{\mathcal{A}}(\alpha') d\alpha' \right\|^2 \\ &= \int \|w(\alpha)\|^2 \rho_{\mathcal{A}}(\alpha) d\alpha - \int \langle w(\alpha), w(\alpha') \rangle \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha' \\ &= \int \left( \|w(\alpha)\|^2 - \langle w(\alpha), w(\alpha') \rangle \right) \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha'. \end{aligned} \quad (40)$$

But, since

$$\int \|w(\alpha)\|^2 \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha = \frac{1}{2} \int (\|w(\alpha)\|^2 + \|w(\alpha')\|^2) \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha', \quad (41)$$

we conclude

$$\int \|w(\alpha) - w_\rho\|^2 \rho_{\mathcal{A}}(\alpha) d\alpha = \frac{1}{2} \int \|w(\alpha) - w(\alpha')\|^2 \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha'. \quad (42)$$

Hence, substituting in (39), we get

$$\text{Var}_\rho(w_\rho)^2 = 2 \int \sigma(\alpha)^2 \rho_{\mathcal{A}}(\alpha) d\alpha + \frac{1}{2} \int \|w(\alpha) - w(\alpha')\|^2 \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha'. \quad (43)$$

**Conditional variance.** We now focus on the conditional variance. As explained in Ex. 1, also in this case, we consider as side information a set of new features  $X = (x_i)_{i=1}^n \in \cup_{n \in \mathbb{N}} \mathcal{X}^n$ . As a consequence, we focus on conditioning functions of the form  $\tau : \cup_{n \in \mathbb{N}} \mathcal{X}^n \rightarrow \mathbb{R}^d$ . From Lemma 2, we know that the ideal function  $\tau_\rho : \cup_{n \in \mathbb{N}} \mathcal{X}^n \rightarrow \mathbb{R}^d$  minimizing the conditional variance term over the space  $\mathcal{T}$  of the measurable functions, is characterized, for almost every  $X \in \cup_{n \in \mathbb{N}} \mathcal{X}^n$ , by

$$\tau_\rho(X) = \mathbb{E}_{\mu \sim \rho(\cdot|X)} w_\mu = \int w_\mu \rho(\mu|X) dw_\mu. \quad (44)$$

We now observe that thanks to the parametrization of the environment  $\rho$ , for any target weight vector  $w_\mu$  and features' set  $X$ , we can write

$$\begin{aligned} \rho(\mu|X) &= \frac{\rho(\mu, X)}{\rho_{\mathcal{X}}(X)} = \frac{\int \rho(\mu, X, \alpha) d\alpha}{\rho_{\mathcal{X}}(X)} = \int \rho(\mu|X, \alpha) \frac{\rho(X, \alpha)}{\rho_{\mathcal{X}}(X)} d\alpha \\ &= \int \rho(\mu|X, \alpha) \rho(\alpha|X) d\alpha = \int \rho(\mu|\alpha) \rho(\alpha|X) d\alpha, \end{aligned} \quad (45)$$

where, in the last equality, we have exploited the fact that, by construction,  $\mu$  is conditionally independent to  $X$  w.r.t.  $\alpha$ , namely  $\rho(\mu|X, \alpha) = \rho(\mu|\alpha)$ . Then, substituting in (44), we get

$$\begin{aligned} \tau_\rho(X) &= \int w_\mu \rho(\mu|X) dw_\mu \\ &= \int w_\mu \left( \int \rho(\mu|\alpha) \rho(\alpha|X) d\alpha \right) dw_\mu \\ &= \int \left( \int w_\mu \rho(\mu|\alpha) dw_\mu \right) \rho(\alpha|X) d\alpha \\ &= \int w(\alpha) \rho(\alpha|X) d\alpha. \end{aligned} \quad (46)$$

**Remark 5** (Asm. 2 in this example). From the expression above, we can conclude that the function  $\tau_\rho$  in (46) is a smooth function of  $X$ , if  $\rho(X|\alpha)$  is a smooth function of  $X$  for any  $\alpha \in \mathcal{A}$ . This means that, in such a case, there exist a Reproducing Kernel Hilbert Space (RKHS)  $\mathcal{H}$  such that  $\tau_\rho \in \mathcal{H}$  and, consequently, making Asm. 2 satisfied. For instance, we can take  $\mathcal{H}$  to be the space induced by the Abel kernel

$$k(X, X') = e^{-\sum_{j=1}^n \frac{\|x_j - x'_j\|}{\sigma}}, \quad \sigma > 0 \quad X, X' \in \cup_{n \in \mathbb{N}} \mathcal{X}^n. \quad (47)$$

In this case,  $\mathcal{H} = W^{d/2+1,2}$  corresponds the Sobolev's space of functions with square integrable  $d/2 + 1$  derivatives.

We now proceed by computing the conditional variance. In order to do this, we observe that

$$\text{Var}_\rho(\tau_\rho)^2 = \mathbb{E}_{(\mu, X) \sim \rho} \|w_\mu - \tau_\rho(X)\|^2 = \mathbb{E}_{X \sim \rho_X} \mathbb{E}_{\mu \sim \rho(\cdot|X)} \|w_\mu - \tau_\rho(X)\|^2. \quad (48)$$

We now observe that, for any set of features  $X$ , exploiting (45), we can rewrite the inner expectation above as follows

$$\begin{aligned} \mathbb{E}_{\mu \sim \rho(\cdot|X)} \|w_\mu - \tau_\rho(X)\|^2 &= \int \|w_\mu - \tau_\rho(X)\|^2 \rho(\mu|X) dw_\mu \\ &= \int \|w_\mu - \tau_\rho(X)\|^2 \left( \int \rho(\mu|\alpha) \rho(\alpha|X) d\alpha \right) dw_\mu \\ &= \int \left( \int \|w_\mu - \tau_\rho(X)\|^2 \rho(\mu|\alpha) dw_\mu \right) \rho(\alpha|X) d\alpha. \end{aligned} \quad (49)$$

But, for each  $\alpha \in \mathcal{A}$ , we can write

$$\begin{aligned} &\int \|w_\mu - \tau_\rho(X)\|^2 \rho(\mu|\alpha) dw_\mu \\ &= \int \|w_\mu\|^2 \rho(\mu|\alpha) dw_\mu - 2 \langle w(\alpha), \tau_\rho(X) \rangle + \|\tau_\rho(X)\|^2 \\ &= \int \|w_\mu\|^2 \rho(\mu|\alpha) dw_\mu \pm \|w(\alpha)\|^2 - 2 \langle w(\alpha), \tau_\rho(X) \rangle + \|\tau_\rho(X)\|^2 \\ &= 2\sigma(\alpha)^2 + \|w(\alpha) - \tau_\rho(X)\|^2. \end{aligned} \quad (50)$$

Hence, substituting into (49), we get

$$\mathbb{E}_{\mu \sim \rho(\cdot|X)} \|w_\mu - \tau_\rho(X)\|^2 = 2 \int \sigma(\alpha)^2 \rho(\alpha|X) d\alpha + \int \|w(\alpha) - \tau_\rho(X)\|^2 \rho(\alpha|X) d\alpha. \quad (51)$$

Hence, integrating w.r.t.  $X$ , we get

$$\begin{aligned} \text{Var}_\rho(\tau_\rho)^2 &= \mathbb{E}_{X \sim \rho_X} \mathbb{E}_{\mu \sim \rho(\cdot|X)} \|w_\mu - \tau_\rho(X)\|^2 \\ &= 2 \int \sigma(\alpha)^2 \rho(\alpha|X) \rho_X(X) d\alpha dX + \int \|w(\alpha) - \tau_\rho(X)\|^2 \rho(\alpha|X) \rho_X(X) d\alpha dX \\ &= 2 \int \sigma(\alpha)^2 \rho_{\mathcal{A}}(\alpha) d\alpha + \int \|w(\alpha) - \tau_\rho(X)\|^2 \rho(\alpha|X) \rho_X(X) d\alpha dX. \end{aligned} \quad (52)$$

We now observe that, exploiting the closed form of  $\tau_\rho$  in (46), the second term above can be rewritten as follows

$$\begin{aligned} &\int \|w(\alpha) - \tau_\rho(X)\|^2 \rho(\alpha|X) \rho_X(X) d\alpha dX \\ &= \int \left\| w(\alpha) - \int w(\alpha') \rho(\alpha'|X) d\alpha' \right\|^2 \rho(\alpha|X) \rho_X(X) d\alpha dX \\ &= \int \|w(\alpha)\|^2 \rho(\alpha|X) \rho_X(X) d\alpha dX - 2 \int \langle w(\alpha), w(\alpha') \rangle \rho(\alpha|X) \rho(\alpha'|X) \rho_X(X) d\alpha d\alpha' dX \\ &\quad + \int \left\| \int w(\alpha') \rho(\alpha'|X) d\alpha' \right\|^2 \rho_X(X) dX. \end{aligned} \quad (53)$$

Note now that

$$\begin{aligned}
& \int \|w(\alpha)\|^2 \rho(\alpha|X) \rho_{\mathcal{X}}(X) d\alpha dX \\
&= \frac{1}{2} \left( \int \|w(\alpha)\|^2 \rho(\alpha|X) \rho_{\mathcal{X}}(X) d\alpha dX + \int \|w(\alpha')\|^2 \rho(\alpha'|X) \rho_{\mathcal{X}}(X) d\alpha' dX \right) \quad (54) \\
&= \frac{1}{2} \left( \int \left( \|w(\alpha)\|^2 + \|w(\alpha')\|^2 \right) \rho(\alpha|X) \rho(\alpha'|X) \rho_{\mathcal{X}}(X) d\alpha d\alpha' dX \right)
\end{aligned}$$

and

$$\int \left\| \int w(\alpha') \rho(\alpha'|X) d\alpha' \right\|^2 \rho_{\mathcal{X}}(X) dX = \int \langle w(\alpha), w(\alpha') \rangle \rho(\alpha|X) \rho(\alpha'|X) \rho_{\mathcal{X}}(X) d\alpha d\alpha' dX. \quad (55)$$

Substituting (54) and (55) in (53), we get

$$\begin{aligned}
& \int \|w(\alpha) - \tau_{\rho}(X)\|^2 \rho(\alpha|X) \rho_{\mathcal{X}}(X) d\alpha dX \\
&= \int \frac{1}{2} \left( \|w(\alpha)\|^2 - 2 \langle w(\alpha), w(\alpha') \rangle + \|w(\alpha')\|^2 \right) \rho(\alpha|X) \rho(\alpha'|X) \rho_{\mathcal{X}}(X) d\alpha d\alpha' dX \\
&= \frac{1}{2} \int \|w(\alpha) - w(\alpha')\|^2 \rho(\alpha|X) \rho(\alpha'|X) \rho_{\mathcal{X}}(X) d\alpha d\alpha' dX \quad (56) \\
&= \frac{1}{2} \int \|w(\alpha) - w(\alpha')\|^2 \frac{\rho(X|\alpha) \rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha' dX \\
&= \frac{1}{2} \int \|w(\alpha) - w(\alpha')\|^2 \left( \int \frac{\rho(X|\alpha) \rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX \right) \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha'.
\end{aligned}$$

Hence, the conditional variance is given by

$$\begin{aligned}
\text{Var}_{\rho}(\tau_{\rho})^2 &= 2 \int \sigma(\alpha)^2 \rho(\alpha) d\alpha \\
&+ \frac{1}{2} \int \|w(\alpha) - w(\alpha')\|^2 \left( \int \frac{\rho(X|\alpha) \rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX \right) \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha'. \quad (57)
\end{aligned}$$

**Conditional vs unconditional variance.** Subtracting (57) to (43), we get that the difference between the unconditional and conditional variance is given by the following closed form

$$\begin{aligned}
& \text{Var}_{\rho}(w_{\rho})^2 - \text{Var}_{\rho}(\tau_{\rho})^2 \\
&= \frac{1}{2} \int \left( 1 - \int \frac{\rho(X|\alpha) \rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX \right) \|w(\alpha) - w(\alpha')\|^2 \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha'. \quad (58)
\end{aligned}$$

Hence, if

$$\int \frac{\rho(X|\alpha) \rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX \leq \epsilon(\alpha, \alpha') \quad (59)$$

for some  $\epsilon : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}_+$ , we can write

$$\text{Var}_{\rho}(w_{\rho})^2 - \text{Var}_{\rho}(\tau_{\rho})^2 \geq \frac{1}{2} \int \left( 1 - \epsilon(\alpha, \alpha') \right) \|w(\alpha) - w(\alpha')\|^2 \rho_{\mathcal{A}}(\alpha) \rho_{\mathcal{A}}(\alpha') d\alpha d\alpha'. \quad (60)$$

## B.2 Clusters (Ex. 1)

The example in the section above encompasses the setting outlined in Ex. 1, by identifying the latent variable  $\alpha$  with the clusters' indexes, namely,  $\mathcal{A} = \{1, \dots, m\}$  and, for any  $\alpha \in \mathcal{A}$ ,  $\rho_{\mathcal{A}}(\alpha) = 1/m$ . We now show that adapting the results above to this specific setting, we manage to show the statement in Ex. 1 in the main body.

**Unconditional variance.** Specifying (43) to the setting outlined in Ex. 1, we get the following closed form for the unconditional variance:

$$\text{Var}_\rho(w_\rho)^2 = \frac{2}{m} \sum_{\alpha=1}^m \sigma(\alpha)^2 + \frac{1}{2m^2} \sum_{\alpha, \alpha'=1}^m \|w(\alpha) - w(\alpha')\|^2. \quad (61)$$

**Conditional variance.** Specifying (57) to the setting outlined in Ex. 1, we get the following closed form for the conditional variance:

$$\text{Var}_\rho(\tau_\rho)^2 = \frac{2}{m} \sum_{\alpha=1}^m \sigma(\alpha)^2 + \frac{1}{2m^2} \sum_{\alpha, \alpha'=1}^m \left( \int \frac{\rho(X|\alpha)\rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX \right) \|w(\alpha) - w(\alpha')\|^2. \quad (62)$$

**Conditional vs unconditional variance.** Finally, specifying (58) to the setting outlined in Ex. 1, we get the following closed form for the gap between the unconditional and the conditional variance:

$$\text{Var}_\rho(w_\rho)^2 - \text{Var}_\rho(\tau_\rho)^2 = \frac{1}{2m^2} \sum_{\alpha, \alpha'=1}^m \left( 1 - \int \frac{\rho(X|\alpha)\rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX \right) \|w(\alpha) - w(\alpha')\|^2. \quad (63)$$

The last ingredient we need to prove the upper bound in Ex. 1 is the following.

**Proposition 6.** *Assume now that for any  $\alpha \in \mathcal{A} = \{1, \dots, m\}$ ,  $\rho(X|\alpha)$  is a Gaussian distribution with mean  $x(\alpha) \in \mathbb{R}^d$  and variance  $\sigma_{\mathcal{X}}^2$ . Then, for any  $\alpha, \alpha' \in \mathcal{A}$ ,*

$$\int \frac{\rho(X|\alpha)\rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX \leq \frac{m}{2} e^{-\frac{n}{\sigma_{\mathcal{X}}^2} \|x(\alpha) - x(\alpha')\|^2}. \quad (64)$$

**Proof.** Thanks to the composition of the environment in clusters, we can write

$$\rho(X) = \sum_{\epsilon=1}^m \rho(X|\epsilon)\rho_{\mathcal{A}}(\epsilon) = \frac{1}{m} \sum_{\epsilon=1}^m \rho(X|\epsilon). \quad (65)$$

As a consequence, for any  $\alpha, \alpha' \in \mathcal{A}$ , we can write

$$\begin{aligned} \int \frac{\rho(X|\alpha)\rho(X|\alpha')}{\rho_{\mathcal{X}}(X)} dX &= m \int \frac{\rho(X|\alpha)\rho(X|\alpha')}{\sum_{\epsilon=1}^m \rho(X|\epsilon)} dX \\ &\leq m \int \frac{\rho(X|\alpha)\rho(X|\alpha')}{\rho(X|\alpha) + \rho(X|\alpha')} dX \\ &\leq \frac{m}{2} \int \sqrt{\rho(X|\alpha)\rho(X|\alpha')} dX, \end{aligned} \quad (66)$$

where in the last inequality we have used the inequality

$$\frac{ab}{a+b} \leq \frac{\sqrt{ab}}{2}, \quad (67)$$

holding for any  $a, b > 0$ . We now observe that, by assumption, we are considering Gaussian distributions for the inputs' probability, i.e., for any  $\alpha \in \{1, \dots, k\}$ , we have

$$\rho(X|\alpha) = \prod_{j=1}^n \frac{1}{\sqrt{2\pi\sigma_{\mathcal{X}}^2}} e^{-\frac{\|x_j - x(\alpha)\|^2}{\sigma_{\mathcal{X}}^2}}. \quad (68)$$

Hence, we have

$$\int \sqrt{\rho(X|\alpha)\rho(X|\alpha')} dX = \frac{1}{\prod_{j=1}^n \sqrt{2\pi\sigma_{\mathcal{X}}^2}} \int e^{-\frac{1}{\sigma_{\mathcal{X}}^2} \sum_{j=1}^n \|x_j - x(\alpha)\|^2 + \|x_j - x(\alpha')\|^2} \prod_j dx_j. \quad (69)$$

We now observe that

$$\begin{aligned}
& \|x_j - x(\alpha)\|^2 + \|x_j - x(\alpha')\|^2 \\
&= 2 \|x_j\|^2 - 2 \langle x_j, x(\alpha) + x(\alpha') \rangle + \|x(\alpha)\|^2 + \|x(\alpha')\|^2 \\
&= 2 \|x_j\|^2 - 2 \langle x_j, x(\alpha) + x(\alpha') \rangle + \|x(\alpha)\|^2 + \|x(\alpha')\|^2 \pm \frac{1}{2} \|x(\alpha) + x(\alpha')\|^2 \\
&= \left\| \sqrt{2}x_j - \frac{1}{\sqrt{2}}(x(\alpha) + x(\alpha')) \right\|^2 - \frac{1}{2} \|x(\alpha) + x(\alpha')\|^2 + \|x(\alpha)\|^2 + \|x(\alpha')\|^2 \\
&= \left\| \sqrt{2}x_j - \frac{1}{\sqrt{2}}(x(\alpha) + x(\alpha')) \right\|^2 + \frac{1}{2} \|x(\alpha)\|^2 + \frac{1}{2} \|x(\alpha')\|^2 - \langle x(\alpha), x(\alpha') \rangle \\
&= \left\| \sqrt{2}x_j - \frac{1}{\sqrt{2}}(x(\alpha) + x(\alpha')) \right\|^2 + \frac{1}{2} \|x(\alpha) - x(\alpha')\|^2.
\end{aligned} \tag{70}$$

Substituting (70) into (69), we conclude

$$\begin{aligned}
& \int \sqrt{\rho(X|\alpha)\rho(X|\alpha')} dX \leq \\
& \leq e^{-\frac{n}{\sigma_{\mathcal{X}}^2} \|x(\alpha) - x(\alpha')\|^2} \frac{1}{\prod_{j=1}^n \sqrt{2\pi\sigma_{\mathcal{X}}^2}} \int e^{-\frac{1}{\sigma_{\mathcal{X}}^2} \sum_{j=1}^n \left\| \sqrt{2}x_j - \frac{1}{\sqrt{2}}(x(\alpha) + x(\alpha')) \right\|^2} \prod_i dx_i \\
& = e^{-\frac{n}{\sigma_{\mathcal{X}}^2} \|x(\alpha) - x(\alpha')\|^2} \frac{1}{\prod_{j=1}^n \sqrt{2\pi\sigma_{\mathcal{X}}^2}} \int e^{-\frac{1}{\sigma_{\mathcal{X}}^2} \sum_{j=1}^n \left\| x_j - \frac{x(\alpha) + x(\alpha')}{2} \right\|^2} \prod_i dx_i \\
& = e^{-\frac{n}{\sigma_{\mathcal{X}}^2} \|x(\alpha) - x(\alpha')\|^2},
\end{aligned} \tag{71}$$

where in the last equality we have exploited the integral of the Gaussian distribution  $\mathcal{N}\left(\frac{x(\alpha) + x(\alpha')}{2}, \sigma_{\mathcal{X}}\right)$ :

$$\frac{1}{\prod_{j=1}^n \sqrt{\pi\sigma_{\mathcal{X}}^2}} \int e^{-\frac{1}{2\sigma_{\mathcal{X}}^2} \sum_{j=1}^n \left\| x_j - \frac{x(\alpha) + x(\alpha')}{2} \right\|^2} \prod_i dx_i = 1. \tag{72}$$

Using the last inequality above in (66), we get the desired statement.  $\blacksquare$

The desired statement in Ex. 1 derives from combining (63) with (64).

### B.3 Circle (Ex. 2)

Consider now the setting outlined in Ex. 2. We proceed as before: we first compute the unconditional variance, then, the conditional variance and, finally, the gap between them.

**Unconditional variance.** We start from observing that, since by construction, for any  $s \in [0, 1]$ ,  $\rho(\mu|s)$  is the Gaussian distribution with mean  $h(s)$ ,  $\rho_S$  is the uniform distribution on  $[0, 1]$  and  $h$  is centered in  $c$ , then, we have

$$w_\rho = \mathbb{E}_{\mu \sim \rho} w_\mu = \mathbb{E}_{s \sim \rho_S} \mathbb{E}_{\mu \sim \rho(\cdot|s)} w_\mu = \mathbb{E}_{s \sim \rho_S} h(s) = c. \tag{73}$$

Hence, we can rewrite the unconditional variance as follows

$$\begin{aligned}
\text{Var}_\rho(w_\rho)^2 &= \mathbb{E}_{(\mu, s) \sim \rho} \|w_\mu - c\|^2 \\
&= \int \|w_\mu - c \pm h(s)\|^2 \rho(\mu, s) dw_\mu ds \\
&= \int \left( \int \|w_\mu - h(s)\|^2 \rho(\mu|s) dw_\mu \right) \rho_S(s) ds + \int \|h(s) - c\|^2 \rho_S(s) ds \\
&\quad + \int \left\langle c - h(s), \int (w_\mu - h(s)) \rho(\mu|s) dw_\mu \right\rangle \rho_S(s) ds \\
&= \sigma^2 + r^2,
\end{aligned} \tag{74}$$

where, in the last equality, we have exploited the fact  $\|h(s) - c\| = r$  for any  $s \in \mathcal{S}$  and the fact that, thanks to the assumption  $\rho(\cdot|s) = \mathcal{N}(h(s), \sigma^2 I)$ ,

$$\int (w_\mu - h(s)) \rho(\mu|s) dw_\mu = 0 \quad \int \|w_\mu - h(s)\|^2 \rho(\mu|s) dw_\mu = \sigma^2. \quad (75)$$

**Conditional variance.** Since by construction  $\rho(\cdot|s) = \mathcal{N}(h(s), \sigma^2 I)$ , we immediately see that the ideal function  $\tau_\rho : [0, 1] \rightarrow \mathbb{R}^d$  in [Lemma 2](#) and the corresponding conditional variance can be, respectively, rewritten as follows

$$\tau_\rho(s) = \mathbb{E}_{\mu \sim \rho(\cdot|s)} w_\mu = h(s) \quad (76)$$

$$\text{Var}_\rho(\tau_\rho)^2 = \mathbb{E}_{(w_\mu, s) \sim \rho} \|w_\mu - h(s)\|^2 = \int \left( \int \|w_\mu - h(s)\|^2 \rho(\mu|s) dw_\mu \right) \rho_S ds = \sigma^2. \quad (77)$$

**Conditional vs unconditional variance.** Subtracting [\(77\)](#) to [\(74\)](#), we get that the difference between the unconditional and conditional variance is given by

$$\text{Var}_\rho(w_\rho)^2 - \text{Var}_\rho(\tau_\rho)^2 = r^2. \quad (78)$$

All the statements given in [Ex. 2](#) have hence been proven.

## C Closed forms for [Asm. 2](#)

Thanks to [Asm. 2](#), we know that there exist  $M_\rho \in \mathbb{R}^{d \times k}$  and  $b_\rho \in \mathbb{R}^d$  such that  $\tau_\rho(\cdot) = M_\rho \Phi(\cdot) + b_\rho$ . In the following lemma, we give the closed form of these quantities and the corresponding variance. We let  $\text{Tr}(\cdot)$  and  $\cdot^*$  be the trace and the conjugate operators respectively.

**Lemma 7** (Best linear conditioning function in hindsight). *Recall the vector  $w_\rho = \mathbb{E}_{\mu \sim \rho_{\mathcal{M}}} w_\mu$  and introduce the vector  $\nu_\rho = \mathbb{E}_{s \sim \rho_S} \Phi(s)$ . Introduce also the following covariance matrices*

$$\text{Cov}_\rho(s, s) = \mathbb{E}_{s \sim \rho_S} \left[ (\Phi(s) - \nu_\rho)(\Phi(s) - \nu_\rho)^\top \right] \in \mathbb{R}^{k \times k} \quad (79)$$

$$\text{Cov}_\rho(w, w) = \mathbb{E}_{\mu \sim \rho_{\mathcal{M}}} \left[ (w_\mu - w_\rho)(w_\mu - w_\rho)^\top \right] \in \mathbb{R}^{d \times d} \quad (80)$$

$$\text{Cov}_\rho(w, s) = \mathbb{E}_{(\mu, s) \sim \rho} \left[ (w_\mu - w_\rho)(\Phi(s) - \nu_\rho)^\top \right] \in \mathbb{R}^{d \times k}. \quad (81)$$

Then,

$$\begin{aligned} \min_{M \in \mathbb{R}^{d \times k}, b \in \mathbb{R}^d} \text{Var}_\rho(\tau_{M,b})^2 &= \text{Var}_\rho(w_\rho)^2 - \text{Tr} \left( \text{Cov}_\rho(w, w) \text{Corr}_\rho(w, s)^\top \text{Corr}_\rho(w, s) \right) \\ &= \text{Var}_\rho(w_\rho)^2 - \left\| \text{Cov}_\rho(w, s) \text{Cov}_\rho(s, s)^{\dagger/2} M_\rho \right\|_F^2, \end{aligned} \quad (82)$$

where we have introduced the correlation matrix

$$\text{Corr}_\rho(w, s) = \text{Cov}_\rho(w, w)^{\dagger/2} \text{Cov}_\rho(w, s) \text{Cov}_\rho(s, s)^{\dagger/2} \in \mathbb{R}^{d \times k}. \quad (83)$$

Moreover, the (minimum norm) values at which the minimum above is attained are given by

$$M_\rho = \text{Cov}_\rho(w, s) \text{Cov}_\rho(s, s)^\dagger \quad (84)$$

$$b_\rho = w_\rho - \text{Cov}_\rho(w, s) \text{Cov}_\rho(s, s)^\dagger \nu_\rho. \quad (85)$$

When [Asm. 2](#) holds, the minimum conditional variance in [Lemma 2](#) can be rewritten as  $\min_{\tau \in \mathcal{T}} \text{Var}_\rho(\tau)^2 = \min_{\tau \in \mathcal{T}_\Phi} \text{Var}_\rho(\tau)^2$ . As a consequence, in this case, the statement above in [\(82\)](#) allows us to express the gap between the conditional and the unconditional variance in [\(12\)](#) as a function of the correlation between the target tasks' weight vectors and the side information. In addition, we can also deduce that such a gap is significant when the 'inclination' of the

linear relation linking the target tasks' weight vectors and the side information (more formally,  $\|\text{Cov}_\rho(s, s)^{1/2} M_\rho\|_F^2$ ) is large. This is not surprising, since, in this case, the gap between conditional and unconditional meta-learning can be interpreted as the gap in using the best linear function w.r.t. the constant one  $\tau \equiv w_\rho$ .

As we will see in the following, the proof of [Lemma 7](#), directly derives from the following facts regarding linear Least Squares.

**Lemma 8.** *Let  $\mathcal{X}$  be an Hilbert space,  $\mathcal{Y} = \mathbb{R}^d$  and  $\mathcal{H} = \mathbb{R}^k$ . Consider a map  $\Psi : \mathcal{X} \rightarrow \mathcal{H}$  and a joint probability distribution  $\rho$  on  $\mathcal{X} \times \mathcal{Y}$  with conditional distribution  $\rho(x|y)$  and marginal  $\rho_Y(y)$ . Denote by  $\otimes$  the standard outer product, introduce the covariance operators:*

$$C_{yy} = \mathbb{E} [y \otimes y] \quad C_{xx} = \mathbb{E} [\Psi(x) \otimes \Psi(x)] \quad C_{xy} = \mathbb{E} [y \otimes \Psi(x)] \quad (86)$$

and the correlation operator

$$\text{Corr}_{xy} = C_{xx}^\dagger C_{xy} C_{yy}^\dagger. \quad (87)$$

Then,

$$\begin{aligned} \min_{M \in \mathbb{R}^{d \times k}} \mathbb{E}_{(x,y) \sim \rho} \|y - M\Psi(x)\|^2 &= \text{Tr}(C_{yy}) - \text{Tr}(C_{yy} \text{Corr}_{xy}^* \text{Corr}_{xy}) \\ &= \text{Tr}(C_{yy}) - \|C_{xx}^{1/2} M_\rho\|_F^2. \end{aligned} \quad (88)$$

The optimal (minimum norm) matrix  $M_\rho$  is given by

$$M_\rho = C_{xy} C_{xx}^\dagger. \quad (89)$$

**Proof.** For any  $M \in \mathbb{R}^{d \times k}$ , we can rewrite

$$\begin{aligned} \mathbb{E}_{(x,y) \sim \rho} \|y - M\Psi(x)\|^2 &= \mathbb{E}_{y \sim \rho_Y} \|y\|^2 + \mathbb{E}_{x \sim \rho_X} \|M\Psi(x)\|^2 - 2 \mathbb{E}_{(x,y) \sim \rho} \langle y, M\Psi(x) \rangle \\ &= \text{Tr}(C_{yy}) + \text{Tr}(C_{xx} M^* M) - 2\text{Tr}(C_{xy}^* M). \end{aligned} \quad (90)$$

By setting the derivatives w.r.t.  $M$  equal to zero, we know that the optimal matrix  $M_\rho$  satisfies

$$M_\rho C_{xx} = C_{xy}. \quad (91)$$

Hence, the optimal (minimum norm) matrix  $M_\rho$  is given by

$$M_\rho = C_{xy} C_{xx}^\dagger. \quad (92)$$

We now compute the corresponding minimum value. We first observe that, by the closed form of the optimal matrix  $M_\rho$ , we can rewrite

$$\begin{aligned} \text{Tr}(C_{xx} M_\rho^* M_\rho) &= \text{Tr}(C_{xx} C_{xx}^\dagger C_{xy}^* C_{xy} C_{xx}^\dagger) = \text{Tr}(C_{xx}^\dagger C_{xx} C_{xx}^\dagger C_{xy}^* C_{xy}) \\ &= \text{Tr}(C_{xx}^\dagger C_{xy}^* C_{xy}), \end{aligned} \quad (93)$$

where in the last equality we have applied the identity  $C_{xx}^\dagger C_{xx} C_{xx}^\dagger = C_{xx}^\dagger$ . We then observe that, again, by the closed form of the optimal matrix  $M_\rho$ , we can rewrite

$$\text{Tr}(C_{xy}^* M_\rho) = \text{Tr}(C_{xy}^* C_{xy} C_{xx}^\dagger) = \text{Tr}(C_{xx}^\dagger C_{xy}^* C_{xy}). \quad (94)$$

Substituting (93) and (94) in (90), we get the following:

$$\min_{M \in \mathbb{R}^{d \times k}} \mathbb{E}_{(x,y) \sim \rho} \|y - M\Psi(x)\|^2 = \text{Tr}(C_{yy}) - \text{Tr}(C_{xx}^\dagger C_{xy}^* C_{xy}) = \text{Tr}(C_{yy}) - \|C_{xx}^{1/2} M_\rho\|_F^2, \quad (95)$$

where in the last equality we have applied the optimality condition (91). In order to terminate the proof, we need to prove the following equality

$$\text{Tr}(C_{xx}^\dagger C_{xy}^* C_{xy}) = \text{Tr}(C_{yy} \text{Corr}_{xy}^* \text{Corr}_{xy}). \quad (96)$$

In order to do this, we proceed as follows. Let  $L^2(\mathcal{Y}, \mathbb{R}, \rho_Y)$  the space of functions from  $\mathcal{Y}$  to  $\mathbb{R}$  that are square integrable w.r.t.  $\rho_Y$  and recall that, for any  $f, g \in L^2(\mathcal{Y}, \mathbb{R}, \rho_Y)$ , such a space is endowed with the scalar product

$$\langle f, g \rangle_{L^2} = \int f(y)g(y) d\rho_Y(y). \quad (97)$$

Throughout the rest of the proof we will use the following operator

$$S : \mathcal{Y} \rightarrow L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}}) \quad h \mapsto (y \mapsto \langle h, \cdot \rangle_{\mathcal{Y}}), \quad (98)$$

where  $\langle \cdot, \cdot \rangle_{\mathcal{Y}}$  is the scalar product in  $\mathcal{Y}$ . Its adjoint operator  $S^* : L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}}) \rightarrow \mathcal{Y}$  is such that, for any  $h \in \mathcal{Y}$  and function  $f \in L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$ ,

$$\langle h, S^* f \rangle_{\mathcal{Y}} = \langle Sh, f \rangle_{L^2} = \int f(y) \langle h, y \rangle_{\mathcal{Y}} d\rho_{\mathcal{Y}}(y) = \left\langle h, \int y f(y) d\rho_{\mathcal{Y}}(y) \right\rangle_{\mathcal{Y}}. \quad (99)$$

This implies that, for any  $f \in L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$ ,

$$S^* f = \int y f(y) d\rho_{\mathcal{Y}}(y). \quad (100)$$

In order to prove the desired statement in (96), we will use the two facts below.

**First fact.** The first fact we need is to show that the operator  $S^* S$  coincides with the covariance operator  $C_{yy}$ , i.e.

$$S^* S = C_{yy} \quad C_{yy} = \mathbb{E} [y \otimes y]. \quad (101)$$

This fact holds, as a matter of fact, we immediately see that, for any  $h_1, h_2 \in \mathcal{Y}$ , we can write

$$\begin{aligned} \langle h_1, S^* S h_2 \rangle_{\mathcal{Y}} &= \langle S h_1, S h_2 \rangle_{L^2} = \int \langle h_1, y \rangle_{\mathcal{Y}} \langle h_2, y \rangle_{\mathcal{Y}} d\rho_{\mathcal{Y}}(y) \\ &= \left\langle h_1, \left( \int y \otimes y d\rho_{\mathcal{Y}}(y) \right) h_2 \right\rangle_{\mathcal{Y}} = \langle h_1, C_{yy} h_2 \rangle_{\mathcal{Y}}. \end{aligned} \quad (102)$$

**Second fact.** Now, recall the map  $\Psi : \mathcal{X} \rightarrow \mathcal{H}$  in the statement and define  $G : \mathcal{Y} \rightarrow \mathcal{H}$  the function

$$G(y) = \int \Psi(x) d\rho(x|y) \quad (103)$$

mapping  $y$  into the conditional expectation of  $\rho(x|y)$ . Assume that  $G \in L^2(\mathcal{Y}, \mathcal{H}, \rho_{\mathcal{Y}})$ , the space of functions from  $\mathcal{Y}$  to  $\mathcal{H}$  that are square integrable w.r.t.  $\rho_{\mathcal{Y}}$ . Note that  $L^2(\mathcal{Y}, \mathcal{H}, \rho_{\mathcal{Y}})$  is isometric to  $\mathcal{H} \otimes L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$ . Denote  $J : L^2(\mathcal{Y}, \mathcal{H}, \rho_{\mathcal{Y}}) \rightarrow \mathcal{H} \otimes L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$  such an isometry and let  $J_G = J(G)$  the Hilbert-Schmidt operator from  $L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$  to  $\mathcal{H}$  associated to  $G$ . Recall that the isometry follows from the observation that, given a basis  $\{h_i\}_{i \in \mathbb{N}}$  of  $\mathcal{H}$  and  $\{f_j\}_{j \in \mathbb{N}}$  of  $L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$ , then the sequence  $\{J(g_{ij})\}_{i,j \in \mathbb{N}}$ , with  $g_{ij}$  the vector-valued functions  $g_{ij}(y) = h_i f_j(y)$  and such that  $J(g_{ij}) = h_i \otimes f_j$ , forms a basis for  $\mathcal{H} \otimes L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$ .

By construction, denoting by  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$  and  $\langle \cdot, \cdot \rangle_{\mathcal{H} \otimes L^2}$  the scalar product in  $\mathcal{H}$  and  $\mathcal{H} \otimes L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$  respectively, for any  $f \in L^2(\mathcal{Y}, \mathbb{R}, \rho_{\mathcal{Y}})$  and  $h \in \mathcal{H}$ , we have

$$\langle J_G, h \otimes f \rangle_{\mathcal{H} \otimes L^2} = \langle G(\cdot), h f(\cdot) \rangle_{L^2(\mathcal{Y}, \mathcal{H}, \rho_{\mathcal{Y}})} = \int \langle G(y), h \rangle_{\mathcal{H}} f(y) d\rho_{\mathcal{Y}}(y). \quad (104)$$

The second fact we need is to show that the operator  $J_G S$  coincides with the covariance operator  $C_{xy}$ , i.e.

$$J_G S = C_{xy} \quad C_{xy} = \mathbb{E} [y \otimes \Psi(x)]. \quad (105)$$

Also this fact holds, as a matter of fact, for any  $h_1 \in \mathcal{H}$  and  $h_2 \in \mathcal{Y}$ , we can write the following

$$\begin{aligned} \langle h_1, J_G S h_2 \rangle_{\mathcal{H}} &= \langle J_G, h_1 \otimes (S h_2) \rangle_{\mathcal{H} \otimes L^2} \\ &= \int \langle G(y), h_1 \rangle_{\mathcal{H}} (S h_2)(y) d\rho_{\mathcal{Y}}(y) \\ &= \int \langle G(y), h_1 \rangle_{\mathcal{H}} \langle h_2, y \rangle_{\mathcal{Y}} d\rho_{\mathcal{Y}}(y) \\ &= \int \langle h_1, (y \otimes G(y)) h_2 \rangle_{\mathcal{H}} d\rho_{\mathcal{Y}}(y) \\ &= \left\langle h_1, \left( \int y \otimes G(y) d\rho_{\mathcal{Y}}(y) \right) h_2 \right\rangle_{\mathcal{H}} \\ &= \langle h_1, C_{xy} h_2 \rangle_{\mathcal{H}}, \end{aligned} \quad (106)$$

where in the last inequality, we have exploited the definition of  $G$  according to which

$$\int y \otimes G(y) d\rho_Y(y) = \int y \otimes \left( \int \Psi(x) d\rho(x|y) \right) d\rho_Y(y) = \int y \otimes \Psi(x) d\rho(x, y) = C_{xy}. \quad (107)$$

As a consequence, recalling the covariance operator  $C_{xx} = \mathbb{E} [\Psi(x) \otimes \Psi(x)]$  and combining the two facts above, we can write the following steps:

$$\begin{aligned} \text{Tr}(C_{xx}^\dagger C_{xy}^* C_{xy}) &= \text{Tr}(C_{xx}^\dagger J_G S S^* J_G^*) \\ &= \text{Tr}(C_{xx}^\dagger J_G S S^\dagger S S^* J_G^*) \\ &= \text{Tr}(C_{xx}^\dagger J_G S S^\dagger S^{*\dagger} S^* S S^* J_G^*) \\ &= \text{Tr}(C_{xx}^\dagger J_G S (S^* S)^\dagger (S^* S) S^* J_G^*) \\ &= \text{Tr}(C_{xx}^\dagger J_G S C_{yy}^\dagger C S^* J_G^*) \\ &= \text{Tr}(C_{yy}^\dagger C_{yy} S^* J_G^* C_{xx}^\dagger J_G S) \\ &= \text{Tr}(C_{yy}^\dagger/2 C_{yy} S^* J_G^* C_{xx}^\dagger J_G S C_{yy}^\dagger/2) \\ &= \text{Tr}(C_{yy} C_{yy}^\dagger/2 S^* J_G^* C_{xx}^\dagger J_G S C_{yy}^\dagger/2) \\ &= \text{Tr}(C_{yy} C_{yy}^\dagger/2 C_{xy} C_{xx}^\dagger C_{xy}^* C_{yy}^\dagger/2) \\ &= \text{Tr}(C_{yy} \text{Corr}_{xy}^* \text{Corr}_{xy}), \end{aligned} \quad (108)$$

where, in the first equation we have used (105), in the second, third and fourth equality we have used the following standard relations

$$S = S S^\dagger S \quad S^\dagger = S^\dagger S^{*\dagger} S^* \quad (S^* S)^\dagger = S^\dagger S^{*\dagger}, \quad (109)$$

in the fifth equality we have used (101), in the eighth equality we have exploited the commuting property  $C_{yy}^\dagger/2 C_{yy} = C_{yy} C_{yy}^\dagger/2$ , in the ninth equality we have used again (105) and, finally, in the last equality, we have introduced the definition of the correlation operator

$$\text{Corr}_{xy} = C_{xx}^\dagger/2 C_{xy} C_{yy}^\dagger/2, \quad (110)$$

which is used in Canonical Correlation Analysis.  $\blacksquare$

We now have all the ingredient for the proof of [Lemma 7](#).

**Proof. of Lemma 7.** We start from recalling the problem we want to solve:

$$\min_{M \in \mathbb{R}^{d \times k}, b \in \mathbb{R}^d} \text{Var}_\rho(\tau_{M,b})^2 = \min_{M \in \mathbb{R}^{d \times k}, b \in \mathbb{R}^d} \mathbb{E}_{(\mu,s) \sim \rho} \|w_\mu - (M\Phi(s) + b)\|^2. \quad (111)$$

By taking the derivatives w.r.t.  $b$ , we conclude that the matrix  $M_\rho \in \mathbb{R}^{d \times k}$  and the vector  $b_\rho \in \mathbb{R}^d$  minimizing the term above satisfy

$$w_\rho = M_\rho \nu_\rho + b_\rho, \quad (112)$$

or, equivalently,

$$b_\rho = w_\rho - M_\rho \nu_\rho. \quad (113)$$

Exploiting this equality, we can rewrite our problem above as

$$\min_{M \in \mathbb{R}^{d \times k}, b \in \mathbb{R}^d} \mathbb{E}_{(\mu,s) \sim \rho} \|w_\mu - (M\Phi(s) + b)\|^2 = \min_{M \in \mathbb{R}^{d \times k}} \mathbb{E}_{(\mu,s) \sim \rho} \|(w_\mu - w_\rho) - M(\Phi(s) - \nu_\rho)\|^2.$$

We now observe that the problem above has the same form of the problem considered in [Lemma 8](#), once one identifies  $\mathcal{X} = \mathcal{S}$  (the space of the side information),  $x = s$ ,  $y = w_\mu - w_\rho$  and  $\Psi(x) = \Phi(s) - \nu_\rho$ . The desired statements automatically derive from the application of [Lemma 8](#) to our context.  $\blacksquare$

## D Proofs of the statements in Sec. 4

In this section we report the proofs of the statements we used in Sec. 4 in order to prove the expected excess risk bound for Alg. 1 in Thm. 4. We start from proving in App. D.1 the properties of the surrogate functions in Prop. 3. Then, in App. D.2, we give the convergence rate of Alg. 1 on the surrogate problem in (18). We conclude by describing in App. D.3 how Alg. 1 can be implemented by computing only evaluations of the kernel associated to the feature map  $\Phi$ , without the need of explicitly evaluating the feature map itself. This is useful when the space in which the image of the feature map lies is high (or even infinite) dimensional.

### D.1 Proof of Prop. 3

We now prove the properties of the surrogate functions in Prop. 3.

**Proposition 3** (Properties of the surrogate meta-loss  $\mathcal{L}$ ). *For any  $Z \in \mathcal{D}$  and  $s \in \mathcal{S}$ , the function  $\mathcal{L}(\cdot, \cdot, s, Z)$  is convex, differentiable and its gradient is given by*

$$\nabla \mathcal{L}(\cdot, \cdot, s, Z)(M, b) = -\lambda \left( A(\tau_{M,b}(s), Z) - \tau_{M,b}(s) \right) \begin{pmatrix} \Phi(s) \\ 1 \end{pmatrix}^\top \quad (19)$$

for any  $M \in \mathbb{R}^{d \times k}$  and  $b \in \mathbb{R}^d$ . Moreover, under Asm. 1 and Asm. 3, we have

$$\|\nabla \mathcal{L}(\cdot, \cdot, s, Z)(M, b)\|_F^2 \leq L^2 R^2 (K^2 + 1). \quad (20)$$

**Proof.** We are interested in studying the properties of the surrogate function

$$\mathcal{L}(\cdot, \cdot, s, Z) : \mathbb{R}^{d \times k} \times \mathbb{R}^d \rightarrow \mathbb{R} \quad (114)$$

in (18). We start from observing that, such a function coincides with the composition of the Moreau envelope  $\hat{\Delta}(\cdot, Z) : \mathbb{R}^d \rightarrow \mathbb{R}$  of the empirical risk  $\mathcal{R}_Z$ :

$$\theta \mapsto \hat{\Delta}(\theta, Z) = \min_{w \in \mathbb{R}^d} \mathcal{R}_{Z,\theta}(w) \quad \mathcal{R}_{Z,\theta}(w) = \frac{1}{n} \sum_{i=1}^n \ell(\langle x_i, w \rangle, y_i) + \frac{\lambda}{2} \|w - \theta\|^2 \quad (115)$$

with the linear transformation

$$s \in \mathcal{S} \mapsto \tau_{M,b}(s) = M\Phi(s) + b \in \mathbb{R}^d. \quad (116)$$

In other words, for any  $M \in \mathbb{R}^{d \times k}$  and  $b \in \mathbb{R}^d$ , we can write

$$\mathcal{L}(M, b, s, Z) = \hat{\Delta}(\tau_{M,b}(s), Z). \quad (117)$$

As a consequence, since the Moreau envelope is convex and differentiable [6, Prop. 12.29], the resulting surrogate function  $\mathcal{L}(\cdot, \cdot, s, Z)$  is convex and differentiable over  $\mathbb{R}^{d \times k} \times \mathbb{R}^d$ . The closed form of the gradient in (19) directly derives from the composition rule for derivatives and the closed form of the gradient of the Moreau envelope [6, Prop. 12.29]

$$\nabla \hat{\Delta}(\cdot, Z)(\theta) = -\lambda (A(\theta, Z) - \theta) \in \mathbb{R}^d, \quad (118)$$

with  $A(\theta, Z)$  defined as in (2). Consequently, we get

$$\nabla \mathcal{L}(\cdot, \cdot, s, Z)(M, b) = \nabla \hat{\Delta}(\cdot, Z)(\tau_{M,b}(s)) \begin{pmatrix} \Phi(s) \\ 1 \end{pmatrix}^\top, \quad (119)$$

coinciding with the desired closed form in (19). Finally, we observe that, as shown in [13, Prop. 4], under Asm. 1, for any  $\theta \in \mathbb{R}^d$ , we have

$$\|\nabla \hat{\Delta}(\cdot, Z)(\theta)\|^2 \leq L^2 R^2. \quad (120)$$

As a consequence, exploiting the rewriting above, Asm. 1 and Asm. 3, we get the desired bound in (20):

$$\begin{aligned} \|\nabla \mathcal{L}(\cdot, \cdot, s, Z)(M, b)\|_F^2 &= \|\nabla \hat{\Delta}(\cdot, Z)(\tau_{M,b}(s)) \Phi(s)^\top\|_F^2 + \|\nabla \hat{\Delta}(\cdot, Z)(\tau_{M,b}(s))\|^2 \\ &= \|\nabla \hat{\Delta}(\cdot, Z)(\tau_{M,b}(s))\|^2 \|\Phi(s)\|^2 + \|\nabla \hat{\Delta}(\cdot, Z)(\tau_{M,b}(s))\|^2 \\ &\leq L^2 R^2 (K^2 + 1), \end{aligned} \quad (121)$$

where in the second equality above we have exploited the fact that for any vectors  $a \in \mathbb{R}^d$  and  $b \in \mathbb{R}^s$ , we have

$$\|ab^\top\|_F^2 = \text{Tr}(ba^\top ab^\top) = \text{Tr}(b^\top ba^\top a) = \|a\|^2 \|b\|^2. \quad (122)$$

■

## D.2 Convergence rate of Alg. 1 on the surrogate problem in (18)

We now give the convergence rate of Alg. 1 on the surrogate problem in (18).

**Proposition 9** (Convergence rate on the surrogate problem in (18)). *Let  $\bar{M}$  and  $\bar{b}$  be the average of the iterations obtained from the application of Alg. 1 over the training data  $(Z_t, s_t)_{t=1}^T$  with constant meta-step size  $\gamma > 0$  and inner regularization parameter  $\lambda > 0$ . Then, under Asm. 1 and Asm. 3, for any  $\tau_{M,b} \in \mathcal{T}_\Phi$ , in expectation w.r.t. the sampling of  $(Z_t, s_t)_{t=1}^T$ ,*

$$\mathbb{E} \hat{\mathcal{E}}_\rho(\tau_{\bar{M}, \bar{b}}) - \hat{\mathcal{E}}_\rho(\tau_{M,b}) \leq \frac{\gamma L^2 R^2 (K^2 + 1)}{2} + \frac{\|(M, b)\|_F^2}{2\gamma T}. \quad (123)$$

**Proof.** We observe that Alg. 1 coincides with Stochastic Gradient Descent applied to the convex and Lipschitz (see Prop. 3) surrogate problem in (18):

$$\min_{M \in \mathbb{R}^{d \times k}, b \in \mathbb{R}^d} \hat{\mathcal{E}}_\rho(\tau_{M,b}) \quad \hat{\mathcal{E}}_\rho(\tau_{M,b}) = \mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^n} \mathcal{L}(M, b, s, Z). \quad (124)$$

As a consequence, by standard arguments (see e.g. [36, Lemma 14.1, Thm. 14.8] and references therein), for any  $\tau_{M,b} \in \mathcal{T}_\Phi$ , we have

$$\mathbb{E} \hat{\mathcal{E}}_\rho(\tau_{\bar{M}, \bar{b}}) - \hat{\mathcal{E}}_\rho(\tau_{M,b}) \leq \frac{\gamma}{2T} \sum_{t=1}^T \mathbb{E} \|\nabla \mathcal{L}(\cdot, \cdot, s, Z_t)(M_t, b_t)\|_F^2 + \frac{\|(M, b)\|_F^2}{2\gamma T}. \quad (125)$$

The desired statement derives from combining this bound with the bound on the norm of the meta-subgradients in (20) in Prop. 3.  $\blacksquare$

## D.3 Implementation of Alg. 1 with kernels

We conclude this section by describing how Alg. 1 can be implemented by computing only evaluations of the kernel associated to the feature map  $\Phi$ . We describe this in the following lemma exploiting standard arguments from online learning with kernels literature (see e.g. [25, 36, 37]).

**Lemma 10** (Implementation of Alg. 1 by kernel's evaluations). *Let  $(M_t, b_t, \theta_t)_{t=1}^T$  be the iteration generated by Alg. 1 with meta-step size  $\gamma \geq 0$ . Then,*

$$\theta_{t+1} = -\gamma \sum_{j=1}^t \nabla \hat{\Delta}(\cdot, Z_j)(\tau_{M_j, b_j}(s_j)) k(s_j, s_{t+1}) + b_{t+1}, \quad (126)$$

where the function  $\hat{\Delta}$  and its gradients  $\nabla \hat{\Delta}(\cdot, Z_j)$  are defined in (115) and (118) above and we have introduced the evaluation

$$k(s_j, s_{t+1}) = \Phi(s_j)^\top \Phi(s_{t+1}), \quad (127)$$

of the kernel associated to the feature map  $\Phi$ .

**Proof.** Exploiting the closed form of the meta-subgradient in (19) in Prop. 3, we can rewrite more explicitly the update step of Alg. 1 as follows:

$$\begin{aligned} M_{t+1} &= M_t - \gamma \nabla \hat{\Delta}(\cdot, Z_t)(\tau_{M_t, b_t}(s_t)) \Phi(s_t)^\top \\ b_{t+1} &= b_t - \gamma \nabla \hat{\Delta}(\cdot, Z_t)(\tau_{M_t, b_t}(s_t)) \\ \theta_{t+1} &= M_{t+1} \Phi(s_{t+1}) + b_{t+1}. \end{aligned} \quad (128)$$

By induction argument on the iteration  $t$ , one can easily see that the update of the matrix  $M_{t+1}$  can be equivalently rewritten as

$$M_{t+1} = -\gamma \sum_{j=1}^t \nabla \hat{\Delta}(\cdot, Z_j)(\tau_{M_j, b_j}(s_j)) \Phi(s_j)^\top. \quad (129)$$

As a consequence, we can rewrite the update of the bias vector  $\theta_{t+1}$  as follows

$$\begin{aligned}
\theta_{t+1} &= M_{t+1}\Phi(s_{t+1}) + b_{t+1} \\
&= -\gamma \sum_{j=1}^t \nabla \hat{\Delta}(\cdot, Z_j)(\tau_{M_j, b_j}(s_j)) \Phi(s_j)^\top \Phi(s_{t+1}) + b_{t+1} \\
&= -\gamma \sum_{j=1}^t \nabla \hat{\Delta}(\cdot, Z_j)(\tau_{M_j, b_j}(s_j)) k(s_j, s_{t+1}) + b_{t+1}.
\end{aligned} \tag{130}$$

This last equation coincides with the desired statement. ■

## E Experimental details

In this section we report the implementation details we omitted in the main body.

In order to tune the hyper-parameters  $\lambda$  and  $\gamma$  our experiments, we followed the same validation procedure described in [13, App. I]. Such a procedure requires performing a meta-training, a meta-validation and a meta-test phase on a separate sets of  $T_{\text{tr}}$  training tasks,  $T_{\text{va}}$  validation tasks and  $T_{\text{te}}$  test tasks. Each task in the training set is observed by a corresponding dataset  $Z_{\text{tr}}$  of  $n = n_{\text{tr}}$  points, while, the tasks in the test and validation sets are all provided with a corresponding training dataset  $Z_{\text{tr}}$  of  $n_{\text{tr}}$  points and a corresponding test dataset  $Z_{\text{te}}$  of  $n_{\text{te}}$  points.

Specifically, in our experiments, we applied the validation procedure above as described in the following.

**Synthetic clusters.** We considered 14 candidates values for both  $\lambda$  and  $\eta$  in the range  $[10^{-5}, 10^5]$  with logarithmic spacing and we evaluated the performance of the estimated meta-parameters (bias vectors) by using  $T = T_{\text{tr}} = 300$ ,  $T_{\text{va}} = 100$ ,  $T_{\text{te}} = 80$  of the available tasks for meta-training, meta-validation and meta-testing, respectively. In order to train and to test the inner algorithm, we splitted each within-task dataset into  $n = n_{\text{tr}} = 50\% n_{\text{tot}}$  for training and  $n_{\text{te}} = 50\% n_{\text{tot}}$  for test.

**Synthetic circle.** We considered 16 candidates values for both  $\lambda$  and  $\eta$  in the range  $[10^{-7}, 10^7]$  with logarithmic spacing and we splitted the data as in the clusters' settings above.

**Lenk dataset.** We considered 14 candidates values for both  $\lambda$  and  $\eta$  in the range  $[10^{-5}, 10^5]$  with logarithmic spacing and we evaluated the performance of the estimated meta-parameters by splitting the tasks into  $T = T_{\text{tr}} = 100$ ,  $T_{\text{va}} = 40$ ,  $T_{\text{te}} = 30$  tasks used for meta-training, meta-validation and meta-testing, respectively. In order to train and to test the inner algorithm, we splitted each within-task dataset into  $n = n_{\text{tr}} = 16$  for training and  $n_{\text{te}} = 4$  for test.

**Schools dataset.** We considered 14 candidates values for both  $\lambda$  and  $\eta$  in the range  $[10^{-5}, 10^5]$  with logarithmic spacing and we evaluated the performance of the estimated meta-parameters by splitting the tasks into  $T = T_{\text{tr}} = 70$ ,  $T_{\text{va}} = 39$ ,  $T_{\text{te}} = 30$  tasks used for meta-training, meta-validation and meta-testing, respectively. In order to train and to test the inner algorithm, we splitted each within-task dataset into  $n = n_{\text{tr}} = 75\% n_{\text{tot}}$  for training and  $n_{\text{te}} = 25\% n_{\text{tot}}$  for test.

We conclude this section reporting the characteristics of the machine we used for running our experiments and the complexity of our method in [Alg. 1](#).

All the experiments were conducted on a workstation with 4 Intel Xeon E5-2697 V3 2.60Ghz CPUs and 256GB RAM.

The variant of our method in [Alg. 1](#) for biased regularization using the batch inner algorithm in (2) has a time and space complexity  $\mathcal{O}(d(k+n))$ . The variant for fine tuning using the online inner algorithm in (3) has a time and space complexity  $\mathcal{O}(dk)$ .