# Conditional Meta-Learning of Linear Representations 

Giulia Denevi ${ }^{1}$, Massimiliano Pontil ${ }^{1,2}$ and Carlo Ciliberto ${ }^{1,2}$<br>${ }^{1}$ University College of London (UK), ${ }^{2}$ Istituto Italiano di Tecnologia (Italy) g.denevi@ucl.ac.uk, massimiliano.pontil@iit.it, c.ciliberto@ucl.ac.uk


#### Abstract

Standard meta-learning for representation learning aims to find a common representation to be shared across multiple tasks. The effectiveness of these methods is often limited when the nuances of the tasks' distribution cannot be captured by a single representation. In this work we overcome this issue by inferring a conditioning function, mapping the tasks' side information (such as the tasks' training dataset itself) into a representation tailored to the task at hand. We study environments in which our conditional strategy outperforms standard meta-learning, such as those in which tasks can be organized in separate clusters according to the representation they share. We then propose a meta-algorithm capable of leveraging this advantage in practice. In the unconditional setting, our method yields a new estimator enjoying faster learning rates and requiring less hyper-parameters to tune than current state-of-the-art methods. Our results are supported by preliminary experiments.


## 1 Introduction

Learning a shared representation among a class of machine learning problems is a wellestablished approach used both in multi-task learning Argyriou et al. (2008); Caruana (1997); Jacob et al. (2009) and meta-learning Balcan et al. (2019); Bertinetto et al. (2018); Bullins et al. (2019); Denevi et al. (2019b); Finn and Levine (2018); Finn et al. (2019); Maurer (2009); Pentina and Lampert (2014); Tripuraneni et al. (2020). The idea behind this methodology is to consider two nested problem: at the within-task level an empirical risk minimization is performed on each task, using inputs transformed by the current representation, on the outer-task (meta-) level, such a representation is updated taking into account the errors of the within-task algorithm on previous tasks.

Such a technique was shown to be advantageous in contrast to solving each task independently when the tasks share a low dimensional representation, see e.g. Balcan et al. (2019); Bullins et al. (2019); Denevi et al. (2019b); Khodak et al. (2019); Maurer (2009); Maurer et al. (2013, 2016); Tripuraneni et al. (2020). However, in real world applications we often deal with heterogeneous classes of learning tasks, which may overall

[^0]be only loosely related. Consequently, the tasks' commonalities may not be captured well by a single representation shared among all the tasks. This is for instance the case in which the tasks can be organized in different groups (clusters), where only tasks belonging to the same cluster share the same low-dimensional representation.

To overcome this issue, in this work, we follow the recent literature on heterogeneous meta-learning Bertinetto et al. (2018); Cai et al. (2020); Denevi et al. (2020); Jerfel et al. (2019); Rusu et al. (2018); Vuorio et al. (2019); Wang et al. (2020); Yao et al. (2019) and propose a so-called conditional meta-learning approach for meta-learning a representation. Our algorithm learns a conditioning function mapping available tasks' side information into a linear representation that is tuned to that task at hand. Our approach borrows from Denevi et al. (2020), where the authors proposed a conditional meta-learning approach for fine tuning and biased regularization. In those cases however, the tasks' target vectors are assumed to be all close to a common bias vector rather than sharing the same lowdimensional linear representation, as instead explored in this work. As we explain in the following, working with linear representations brings additional difficulties than working with bias vectors, but, on the other hand, it is also a relevant and effective framework in many scenarios.

In this work, we propose an online conditional method for linear representation learning with strong theoretical guarantees. In particular, we show that the method is advantageous over standard (unconditional) representation learning methods used in meta-learning when the environment of observed tasks is heterogeneous.

Contributions and Organization. The contributions of this work are the following. First, in Sec. 2, we design a conditional meta-learning approach to infer a linear representation that is tuned to the task at hand. Second, in Sec. 3, we formally characterize circumstances under which our conditional framework brings advantage w.r.t. the standard unconditional approach. In particular, we argue that this is the case when the tasks are organized in different clusters according to the support pattern or linear representation their target vectors' share. Third, in Sec. 4, we design a convex meta-algorithm providing a comparable gain as the number of the tasks it observes increases. In the unconditional setting, the proposed method is able to recover faster rates and it requires to tune one less hyperparameter w.r.t. the state-of-the-art unconditional methods. Finally, in Sec. 5, we present numerical experiments supporting our theoretical claims. We conclude our work in Sec. 6 and we postpone the missing proofs to the supplementary material.

## 2 Conditional Representation Learning

In this section we introduce our conditional meta-learning setting for representation learning. Then, we proceed to identify the differences w.r.t. (with respect to) the standard unconditional counterpart. We begin our overview by first introducing the class of inner learning algorithms considered in this work.

Within-Task Algorithms. We consider the standard linear supervised learning setting over $\mathcal{Z}=\mathcal{X} \times \mathcal{Y}$ with $\mathcal{X} \subseteq \mathbb{R}^{\mathrm{d}}$ and $\mathcal{Y} \subseteq \mathbb{R}$ input and output spaces, respectively. We denote by $\mathcal{P}(\mathcal{Z})$ the set of probability distributions (tasks) over $\mathcal{Z}$. For any task $\mu \in \mathcal{P}(\mathcal{Z})$ and a given
loss function $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, we aim at finding a weight vector $w_{\mu} \in \mathbb{R}^{d}$ minimizing the expected risk

$$
\begin{equation*}
\min _{w \in \mathbb{R}^{\mathrm{d}}} \mathcal{R}_{\mu}(w) \quad \mathcal{R}_{\mu}(w)=\mathbb{E}_{(x, y) \sim \mu} \ell(\langle x, w\rangle, y), \tag{1}
\end{equation*}
$$

where, $\langle\cdot, \cdot\rangle$ represents the Euclidean product in $\mathbb{R}^{\text {d }}$. In practice, $\mu$ is only partially observed trough a dataset $Z=\left(x_{i}, y_{i}\right)_{i=1}^{n} \sim \mu^{n}$, namely, a collection of $n$ identically independently distributed (i.i.d.) points sampled from $\mu$. Thus, the goal becomes to use a learning algorithm in order to estimate a candidate weight vector with a small expected risk converging to the ideal $\mathcal{R}_{\mu}\left(w_{\mu}\right)$ as the sample size $n$ grows.

Specifically, in this work we will consider as candidate estimators, the family of regularized empirical risk minimizers for linear feature learning Argyriou et al. (2008). Formally, denoting by $\mathcal{D}=\bigcup_{n \in \mathbb{N}} \mathcal{Z}^{n}$ the space of all datasets on $\mathcal{Z}$, for a given $\theta \in \Theta$ in $\Theta=\mathbb{S}_{+}^{d}$ the set of positive definite $\mathrm{d} \times \mathrm{d}$ matrices, we will consider the following learning algorithms $A(\theta, \cdot): \mathcal{D} \rightarrow \mathbb{R}^{\mathrm{d}}:$

$$
\begin{equation*}
\mathcal{A}(\theta, Z)=\underset{w \in \operatorname{Ran}(\theta) \subset \mathbb{R}^{\mathbf{d}}}{\operatorname{argmin}} \mathcal{R}_{Z, \theta}(w), \tag{2}
\end{equation*}
$$

where $\operatorname{Ran}(\theta)$ denotes the range of $\theta$ and we defined

$$
\begin{equation*}
\mathcal{R}_{z, \theta}(w)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\left\langle x_{i}, w\right\rangle, y_{i}\right)+\frac{1}{2}\left\langle w, \theta^{\dagger} w\right\rangle, \tag{3}
\end{equation*}
$$

for any $w \in \operatorname{Ran}(\theta)$. Here $\theta^{\dagger}$ denotes the pseudoinverse of $\theta$. Throughout this work we will denote by $\mathcal{R}_{Z}(\cdot)=1 / n \sum_{i=1}^{n} \ell\left(\left\langle x_{i}, \cdot \cdot\right\rangle, y_{i}\right)$ the empirical risk associated to $Z$.

Here, $\theta$ plays the role of a linear feature representation that is learned during the meta-learning process (see Argyriou et al., 2008, for more details on the interpretation).

Remark 1 (Within-Task Regularization Parameter). We observe that, differently to previous work (see e.g. Denevi et al., 2019b), we consider the meta-parameters $\theta$ to be any positive semidefinite matrix, without constraint on its trace (e.g. $\operatorname{Tr}(\theta) \leq 1$ ). This allows us to absorb the regularization parameter $\lambda$ typically used to control $\lambda\left\langle w, \theta^{\dagger} w\right\rangle$. This choice is advantageous both in practice since it reduces the number of hyper-parameter to tune and in theory (as discussed in the following) by enjoying faster learning rates.

Remark 2 (Online Variant of Eq. (2)). While in the following we will focus on algorithms of the form of Eq. (2), our analysis and results extend also to the setting in which the exact minimization of the empirical risk is replaced by a pre-conditioned variant of online gradient descent on $\mathcal{R}_{Z, \theta}$, with starting point $w_{0}=0 \in \mathbb{R}^{\mathrm{d}}$ and step size inversely proportional to the iteration:

$$
\begin{align*}
& A(\theta, Z)=\frac{1}{n} \sum_{i=1}^{n} w_{i}, \quad w_{i+1}=w_{i}-\frac{\theta p_{i}}{i}  \tag{4}\\
& p_{i}=s_{i} x_{i}+\theta^{\dagger} w_{i} \quad s_{i} \in \partial \ell\left(\cdot, y_{i}\right)\left(\left\langle x_{i}, w_{i}\right\rangle\right) .
\end{align*}
$$

This modification brings additional negligible logarithmic factors in our bounds in the following.

Unconditional Meta-Learning. The standard unconditional meta-learning setting assumes there exist a meta-distribution $\rho \in \mathcal{P}(\mathcal{M})$ - also called environment in (Baxter, 2000) - over a family $\mathcal{M} \subseteq \mathcal{P}(\mathcal{Z})$ of distributions (tasks) $\mu$ and it aims at selecting an inner algorithm in the family above that is well suited to solve tasks $\mu$ sampled from $\rho$. This target can be reformulated as finding a linear representation $\theta_{\rho} \in \Theta$ such that the corresponding algorithm $\mathcal{A}\left(\theta_{\rho}, \cdot\right)$ minimizes the transfer risk

$$
\begin{equation*}
\min _{\theta \in \Theta} \mathcal{E}_{\rho}(\theta) \quad \mathcal{E}_{\rho}(\theta)=\mathbb{E}_{\mu \sim \rho} \mathbb{E}_{Z \sim \mu^{n}} \mathcal{R}_{\mu}(A(\theta, Z)) . \tag{5}
\end{equation*}
$$

In practice, this stochastic problem is usually tackled by iteratively sampling a task $\mu \sim \rho$ and a corresponding dataset $Z \sim \mu^{n}$, and, then, performing a step of stochastic gradient descent on an empirical approximation of Eq. (5) computed from Z. This has approach has proven effective for instance when the tasks of the environment share a simple common linear representation, see e.g. Balcan et al. (2019); Bullins et al. (2019); Denevi et al. (2019a,b); Finn et al. (2017); Finn and Levine (2018); Finn et al. (2019); Khodak et al. (2019). However, when a single linear representation is not sufficient for the entire environment of tasks (e.g. multi-clusters), this homogeneous approach is expected to fail. In order to overcome this limitation, some recent works have adopted the following conditional approach to the problem, see e.g. Cai et al. (2020); Denevi et al. (2020); Jerfel et al. (2019); Rusu et al. (2018); Vuorio et al. (2019); Wang et al. (2020); Yao et al. (2019).

Conditional Meta-Learning. Analogously to Denevi et al. (2020), we assume that any task $\mu \sim \rho$ is provided of additional side information $s \in \mathcal{S}$. In such a case, we consider the environment $\rho$ as 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. Moreover, as usual, we assume $\rho$ to decompose in $\rho(\cdot \mid s) \rho_{\mathcal{S}}(\cdot)$ and $\rho(\cdot \mid \mu) \rho_{\mathcal{M}}(\cdot)$ the conditional and marginal distributions w.r.t. $\mathcal{S}$ and $\mathcal{M}$. For instance, we observe that the side information $s$ could contain descriptive features of the associated task, for example attributes in collaborative filtering Abernethy et al. (2009), or additional information about the users in recommendation systems Harper and Konstan (2015)). Moreover $s$ could be formed by a portion of the dataset sampled from $\mu$ (see Denevi et al. (2020); Wang et al. (2020)). Conditional meta-learning leverages this additional side information in order to adapt (or condition) the linear representation $\theta \in \Theta$ on the associated task at hand, by learning a linear-representation-valued function $\tau$ solving the problem

$$
\begin{equation*}
\min _{\tau \in \mathcal{T}} \mathcal{E}_{\rho}(\tau), \quad \mathcal{E}_{\rho}(\tau)=\mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{\mathrm{Z} \sim \mu^{n}} \mathcal{R}_{\mu}(A(\tau(\mathrm{~s}), \mathrm{Z})) \tag{6}
\end{equation*}
$$

over the space $\mathcal{T}$ of measurable functions $\tau: \mathcal{S} \rightarrow \Theta$. Notice that we retrieve the unconditional meta-learning problem in Eq. (5) if we restrict Eq. (6) to the set of functions $\mathcal{T}^{\text {const }}=\{\tau \mid \tau(\cdot) \equiv \theta, \theta \in \Theta\}$, mapping all the side information into the same constant linear representation.

In the next section, we will investigate the theoretical advantages of adopting such a conditional perspective and, then, we will introduce a convex meta-algorithm to tackle Eq. (6).

## 3 The Advantage of Conditional Representation Learning

In order to characterize the behavior of the optimal solution of Eq. (6) and to investigate the potential advantage of conditional meta-learning, we analyze the generalization properties of a given conditioning function $\tau$. Formally, we compare the $\operatorname{error} \mathcal{E}_{\rho}(\tau)$ w.r.t. the optimal minimum risk

$$
\begin{equation*}
\mathcal{E}_{\rho}^{*}=\mathbb{E}_{\mu \sim \rho} \mathcal{R}_{\mu}\left(w_{\mu}\right) \quad w_{\mu}=\underset{w \in \mathbb{R}^{d}}{\operatorname{argmin}} \mathcal{R}_{\mu}(w) . \tag{7}
\end{equation*}
$$

In order to do this, we first need to introduce the following standard assumptions used also in previous literature. Throughout this work we will denote by ${ }^{\top}$ the standard transposition operation.

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. For any $s \sim \rho_{\mathcal{S}}$, introduce the conditional covariance matrices

$$
\begin{equation*}
W(s)=\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} w_{\mu} w_{\mu}^{\top}, \quad C(s)=\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} \mathbb{E}_{x \sim \eta_{\mu}} x x^{\top} \tag{8}
\end{equation*}
$$

where, $\eta_{\mu}$ denotes the inputs' marginal distribution of the task $\mu$. Let $\tau \in \mathcal{T}$ such that $\operatorname{Ran}(\mathcal{W}(s)) \subseteq \operatorname{Ran}(\tau(s))$ for any $s \sim \rho_{\mathcal{S}}$ and let $\mathcal{A}(\tau(s), \cdot)$ be the associated inner algorithm from Eq. (2). Then,

$$
\begin{equation*}
\mathcal{E}_{\rho}(\tau)-\mathcal{E}_{\rho}^{*} \leq \frac{\mathbb{E}_{s \sim \rho_{\mathcal{S}}} \operatorname{Tr}\left(\tau(s)^{\dagger} W(s)\right)}{2}+\frac{2 L^{2} \mathbb{E}_{s \sim \rho_{\mathcal{S}}} \operatorname{Tr}(\tau(s) \mathrm{C}(s))}{n} . \tag{9}
\end{equation*}
$$

Proof. For any $(\mu, s) \sim \rho$, consider the decomposition

$$
\begin{equation*}
\mathcal{E}_{\rho}(\tau)-\mathcal{E}_{\rho}^{*}=\mathbb{E}_{(\mu, \mathrm{s}) \sim \rho}\left[\mathrm{B}_{\mu, \mathrm{s}}+\mathrm{C}_{\mu, \mathrm{s}}\right], \tag{10}
\end{equation*}
$$

with
$B_{\mu, s}=\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{\mu}(A(\tau(s), Z))-\mathcal{R}_{Z}(A(\tau(s), Z))\right] C_{\mu, s}=\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z}(A(\tau(s), Z))-\mathcal{R}_{\mu}\left(w_{\mu}\right)\right]$.
$B_{\mu, s}$ is the generalization error of the inner algorithm $A(\tau(s), \cdot)$ on the task $\mu$. Hence, applying stability arguments (see Prop. 6 in App. A), we can write

$$
\mathrm{B}_{\mu, s} \leq \frac{2 \mathrm{~L}^{2} \operatorname{Tr}\left(\tau(\mathrm{~s}) \mathbb{E}_{x \sim \eta_{\mu}} \chi \chi^{\top}\right)}{n}
$$

Regarding the term $\mathrm{C}_{\mu, s}$, for any conditioning function $\tau$ such that $w_{\mu} \in \operatorname{Ran}(\tau(s))$, we can write

$$
\begin{aligned}
\mathcal{C}_{\mu, s} & =\mathbb{E}_{Z \sim \mu^{n}}\left[\min _{w \in \mathbb{R}^{:}: w \in \operatorname{Ran}(\tau(s))} \mathcal{R}_{Z, \tau(s)}(w)-\mathcal{R}_{\mu}\left(w_{\mu}\right)\right] \\
& \leq \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}\left(w_{\mu}\right)-\mathcal{R}_{\mu}\left(w_{\mu}\right)\right] \\
& =\frac{\operatorname{Tr}\left(\tau(s)^{\dagger} w_{\mu} w_{\mu}^{\top}\right)}{2},
\end{aligned}
$$

where, the second equality exploits the definition of the algorithm in Eq. (2) and the first inequality exploits the definition of minimum. The desired statement follows by combining the two bounds above, rewriting $\mathbb{E}_{(\mu, s) \sim \rho}=\mathbb{E}_{s \sim \rho_{s}} \mathbb{E}_{\mu \sim \rho(\cdot \mid s)}$ and observing that the constraint above on $\tau$ can be rewritten as follows

$$
\begin{aligned}
w_{\mu} \in \operatorname{Ran}(\tau(s)) \text { for any }(\mu, s) \sim \rho & \Longleftrightarrow \operatorname{Ran}\left(w_{\mu} w_{\mu}^{\top}\right) \subseteq \operatorname{Ran}(\tau(s)) \text { for any }(\mu, s) \sim \rho \\
& \Longleftrightarrow \mathbb{E}_{\mu \sim p(\cdot \mid s)}\left[\operatorname{Ran}\left(w_{\mu} w_{\mu}^{\top}\right)\right] \subseteq \operatorname{Ran}(\tau(s)) \text { for any } s \sim \rho_{\mathcal{S}} \\
& \Longleftrightarrow \operatorname{Ran}\left(\mathbb{E}_{\mu \sim \rho(\cdot \mid s)}\left[w_{\mu} w_{\mu}^{\top}\right]\right) \subseteq \operatorname{Ran}(\tau(s)) \text { for any } s \sim \rho_{\mathcal{S}},
\end{aligned}
$$

where the second and the third equivalences derive from the fact that, for any matrices $A, B \in \mathbb{S}_{+}^{d}$ and any scalar $c \neq 0, \operatorname{Ran}(\mathcal{A}) \subseteq \operatorname{Ran}(A+B)=\operatorname{Ran}(A)+\operatorname{Ran}(B)$ and $\operatorname{Ran}(c \mathcal{A})=$ $\operatorname{Ran}(A)$, see e.g. Hogben $(2006,2013)$.

Thm. 1 suggests that the conditioning function $\tau_{*}$ minimizing the right hand side of Eq. (9) is a good candidate to solve the meta-learning problem. The following result explores this question by showing that such a minimizer admits a closed form solution. The proof is reported in App. B. In the following, we will denote by $\|\cdot\|_{F}$ and $\|\cdot\|_{*}$ the Frobenius and trace norm of a matrix, respectively.

Proposition 2 (Best Conditioning Function in Hindsight). The conditioning function minimizer and the minimum of the bound presented in Thm. 1 over the set

$$
\left\{\tau \in \mathcal{T} \mid \operatorname{Ran}(W(s)) \subseteq \operatorname{Ran}(\tau(s)), \rho_{\mathcal{S}} \text {-almost surely }\right\},
$$

are respectively

$$
\tau_{\rho}(s)=\frac{\sqrt{n}}{2 L} C(s)^{\dagger / 2}\left(C(s)^{1 / 2} W(s) C(s)^{1 / 2}\right)^{1 / 2} C(s)^{\dagger / 2}
$$

and

$$
\begin{equation*}
\mathcal{E}_{\rho}\left(\tau_{\rho}\right)-\mathcal{E}_{\rho}^{*} \leq \frac{2 \mathrm{LE}_{s \sim \rho \mathcal{S}}\left\|\mathrm{~W}(s)^{1 / 2} \mathrm{C}(s)^{1 / 2}\right\|_{*}}{\sqrt{n}} \tag{11}
\end{equation*}
$$

This result allows us to quantify the benefits of adopting the conditional feature learning strategy.

Conditional Vs. Unconditional Meta-Learning. Applying Prop. 2 to $\mathcal{T}^{\text {const }}$, we obtain the excess risk bound for unconditional meta-learning

$$
\begin{equation*}
\mathcal{E}_{\rho}\left(\theta_{\rho}\right)-\mathcal{E}_{\rho}^{*} \leq \frac{2 L\left\|W_{\rho}^{1 / 2} C_{\rho}^{1 / 2}\right\|_{*}}{\sqrt{n}} \tag{12}
\end{equation*}
$$

achieved for $\tau(s) \equiv \theta_{\rho}$ the meta-parameter

$$
\begin{equation*}
\theta_{\rho}=\frac{\sqrt{n}}{2 L} C_{\rho}^{\dagger / 2}\left(C_{\rho}^{1 / 2} W_{\rho} C_{\rho}^{1 / 2}\right)^{1 / 2} C_{\rho}^{\dagger / 2}, \tag{13}
\end{equation*}
$$

with unconditional covariance matrices

$$
\begin{equation*}
W_{\rho}=\mathbb{E}_{\mu \sim \rho} w_{\mu} w_{\mu}^{\top}, \quad C_{\rho}=\mathbb{E}_{\mu \sim \rho} \mathbb{E}_{x \sim \eta_{\mu}} x x^{\top} . \tag{14}
\end{equation*}
$$

We observe that in the previous literature Denevi et al. $(2018,2019 b)$ the authors restricted the unconditional problem over the smaller class of linear representation $\widehat{\Theta}=$ $\left\{\theta \in \mathbb{S}_{+}^{d}: \operatorname{Ran}\left(W_{\rho}\right) \subseteq \operatorname{Ran}(\theta), \operatorname{Tr}(\theta) \leq 1\right\}$ and they considered as the best unconditional representation, the matrix minimizing only a part of the previous bound, namely,

$$
\begin{equation*}
\hat{\theta}_{\rho}=\underset{\theta \in \Theta}{\operatorname{argmin}} \operatorname{Tr}\left(\theta^{\dagger} W_{\rho}\right)=\frac{W_{\rho}^{1 / 2}}{\operatorname{Tr}\left(W_{\rho}^{1 / 2}\right)} . \tag{15}
\end{equation*}
$$

On the other hand, the unconditional oracle we introduce above in Eq. (13) allows us to recover a tighter bound which is able to recover the best performance between independent task learning (ITL) and the oracle considered in previous literature Denevi et al. (2019b). Indeed, by exploiting the duality between the trace norm $\|\cdot\|_{*}$ and the operator norm $\|\cdot\|_{\infty}$ of a matrix, we can upper bound the right-side-term in Eq. (12) by the quantity

$$
\frac{2 L \min \left\{\left\|W_{\rho}^{1 / 2}\right\|_{*}\left\|C_{\rho}^{1 / 2}\right\|_{\infty},\left\|W_{\rho}^{1 / 2}\right\|_{F}\left\|C_{\rho}^{1 / 2}\right\|_{F}\right\}}{\sqrt{n}},
$$

namely, the minimum between the bound for independent task learning and the bound for unconditional oracle obtained by previous authors. Notice that the unconditional quantity in Eq. (12) is always bigger than the conditional quantity in Eq. (11), since Eq. (12) coincides with the minimum over a smaller class of function. In order to quantify the gap between these two quantities - namely, the advantage in using the conditional approach w.r.t. the unconditional one - we have to compare the term $\left\|W_{\rho}^{1 / 2} C_{\rho}^{1 / 2}\right\|_{*}$ with the term $\mathbb{E}_{s \sim \rho_{\mathcal{S}}}\left\|C(s)^{1 / 2} W(s)^{1 / 2}\right\|_{*}$.

We report below a setting that can be considered illustrative for many real-world scenarios in which such a gap in performance is significant. We refer to App. C for the details and the deduction.

Example 1 (Clusters). Let $\mathcal{S}=\mathbb{R}^{q}$ be the side information space, for some integer $q>0$. Let $\rho$ be such that the side information marginal distribution $\rho_{\mathcal{S}}$ is given by a uniform mixture of $\mathfrak{m}$ uniform distributions. More precisely, let $\rho_{\mathcal{S}}=\frac{1}{m} \sum_{i=1}^{m} \rho_{\mathcal{S}}^{(i)}$, with $\rho_{\mathcal{S}}^{(i)}=\mathcal{U}\left(\mathcal{B}\left(a_{i}, 1 / 2\right)\right)$ the uniform distribution on the ball of radius $1 / 2$ centered at $a_{i} \in \mathcal{S}$, characterizing the cluster i. For a given side information s, a task $\mu \sim \rho(\cdot \mid s)$ is sampled such that: 1 ) its inputs' marginal $\eta_{\mu}$ is a distribution with constant covariance matrix $C(s)=\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} \mathbb{E}_{x \sim \eta_{\mu}} \chi x^{\top}=C$, for some $\mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}}$, 2) $w_{\mu}$ is sampled from a distribution with conditional covariance matrix $W(s)=\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} w_{\mu} w_{\mu}^{\top}$, with $W(s)$ such that $\left(\mathrm{C}^{1 / 2} \mathrm{~W}(\mathrm{~s}) \mathrm{C}^{1 / 2}\right)\left(\mathrm{C}^{1 / 2} \mathrm{~W}(\mathrm{p}) \mathrm{C}^{1 / 2}\right)=0$ if $\mathrm{s} \neq \mathrm{p}$. Then,

$$
\mathbb{E}_{s \sim \rho_{\mathcal{S}}}\left\|C(s)^{1 / 2} W(s)^{1 / 2}\right\|_{*}=\frac{1}{\sqrt{m}}\left\|W_{\rho}^{1 / 2} C_{\rho}^{1 / 2}\right\|_{*} .
$$

The inequality above tells us that, in the setting of Ex. 1, the conditional approach gains a $\sqrt{m}$ factor in comparison to the unconditional approach. Therefore, the larger the number of clusters is, the more pronounced the advantage of conditional approach w.r.t. the unconditional one will be. We observe that a particular case of the setting above could be that one in which $\mathrm{q}=1$ and the side information are noisy observations of the index of the cluster the tasks belong to. In our experiments, in Sec. 5, we consider a more
interesting and realistic variant of the setting above, in which we will use as task's side information a training dataset sampled from that task. In the next section, we introduce a convex meta-algorithm mimicking this advantage also in practice.

## 4 Conditional Representation Meta-Learning Algorithm

To tackle conditional meta-learning in practice we consider a parametrization where the conditioning functions that are modeled w.r.t. a given feature map $\Phi: \mathcal{S} \rightarrow \mathbb{R}^{k}$ (with $k \in \mathbb{N}$ ) on the side information space. In other words, we consider $\tau: \mathcal{S} \rightarrow \mathbb{S}_{+}^{d}$,

$$
\begin{equation*}
\tau(\cdot)=(M \Phi(\cdot))^{\top} M \Phi(\cdot)+C, \tag{16}
\end{equation*}
$$

for some tensor $M \in \mathbb{R}^{p \times d \times k}(p \in \mathbb{N})$ and matrix $C \in \mathbb{S}_{+}^{d}$.
By construction, the above parametrization guarantees us to learn functions taking values in the set of positive semi-definite matrices. However, directly addressing the metalearning problem poses two issues: first, dealing with tensorial structures might become computationally challenging in practice and second, such parametrization is quadratic in $M$ and would lead to a non-convex optimization functional in practice. To tackle this issue, the following results shows that we can equivalently rewrite the conditioning function in the form of Eq. (16) by using a matrix in $\mathbb{S}_{+}^{\mathrm{dk}}$. This will allows us to implement our method working with matrices in $\mathbb{S}_{+}^{\text {dk }}$, instead of tensors in $\mathbb{R}^{p \times d \times k}$. Throughout this work, we will denote by $\otimes$ the Kronecker product.

Proposition 3 (Matricial Re-formulation of $\tau_{M}(s)$ ). Let $\tau$ be as in Eq. (16). Then,

$$
\begin{equation*}
\tau(s)=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)^{\top}\right) \mathrm{H}_{M}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right)+\mathrm{C}, \tag{17}
\end{equation*}
$$

where $\mathrm{I}_{\mathrm{d}}$ is the identity in $\mathbb{R}^{\mathrm{d} \times \mathrm{d}}$ and $\mathrm{H}_{\mathrm{M}}$ is the matrix in $\mathbb{R}^{\mathrm{dk} \times \mathrm{dk}}$ defined by the entries

$$
\left(\mathrm{H}_{\mathrm{M}}\right)_{(\mathrm{i}-1) \mathrm{k}+\mathrm{h},(\mathrm{j}-1) \mathrm{k}+z}=\langle\mathrm{M}(:, \mathfrak{i}, \mathrm{h}), \mathrm{M}(:, \mathfrak{j}, z)\rangle
$$

with $i, j=1, \ldots, d$ and $h, z=1, \ldots, k$.
The arguments above motivate us to consider the following set of conditioning functions:

$$
\begin{equation*}
\mathcal{T}_{\Phi}=\left\{\tau(\cdot)=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\cdot)^{\top}\right) \mathrm{H}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\cdot)\right)+\mathrm{C} \mid \text { such that } \mathrm{H} \in \mathbb{S}_{+}^{\mathrm{dk}}, \mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}}\right\} . \tag{18}
\end{equation*}
$$

To highlight the dependency of a function $\tau \in \mathcal{T}_{\Phi}$ w.r.t. its parameter H and C , we will denote $\tau=\tau_{\mu, C}$. 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 $\left\|\tau_{\mathrm{H}, \mathrm{C}}\right\|^{2}=\|(\mathrm{H}, \mathrm{C})\|_{\mathrm{F}}^{2}=\|\mathrm{H}\|_{\mathrm{F}}^{2}+\|\mathrm{C}\|_{\mathrm{F}}^{2}$, where, recall, $\|\cdot\|_{F}$ denotes the Frobenius norm. The following two standard assumptions will allow us to design and analyse our method.

Assumption 2. The optimal function $\tau_{\rho}$ belongs to $\mathcal{T}_{\Phi}$, namely there exist $\mathrm{H}_{\rho} \in \mathbb{S}_{+}^{\mathrm{dk}}$ and $\mathrm{C}_{\rho} \in \mathbb{S}_{+}^{\mathrm{d}}$, such that $\tau_{\rho}(\cdot)=\tau_{\mathrm{H}_{\rho}, \mathrm{C}_{\rho}}(\cdot)=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\cdot)^{\top}\right) \mathrm{H}_{\rho}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\cdot)\right)+\mathrm{C}_{\rho}$.
Assumption 3. There exists $K>0$ such that $\|\Phi(s)\| \leq K$ for any $s \in \mathcal{S}$.

Here, Asm. 2 allows us to restrict the conditional meta-learning problem in Eq. (6) to $\mathcal{T}_{\Phi}$, rather than to the entire space $\mathcal{T}$ of measurable functions, while Asm. 3 ensures that the meta-objective is Lipschitz (see below).

The Convex Surrogate Problem. We start from observing that, exploiting the generalization properties of the within-task algorithm (see Prop. 6 in App. A), for any $\tau$, we can write the following

$$
\begin{aligned}
\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{\mu}(A(\tau(s), Z))\right] \leq & \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z}(A(\tau(s), Z))\right]+\frac{2 L^{2} \operatorname{Tr}\left(\tau(s) \mathbb{E}_{\chi \sim \eta_{\mu}} \chi x^{\top}\right)}{n} \\
& \leq \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}(A(\tau(s), Z))\right]+\frac{2 L^{2} \operatorname{Tr}\left(\tau(s) \mathbb{E}_{\chi \sim \eta_{\mu}} x x^{\top}\right)}{n},
\end{aligned}
$$

where in the second inequality we have exploited the fact that the within-task regularizer is non-negative. Consequently, by taking the expectation w.r.t. $(\mu, s) \sim \rho$ and exploiting the fact that the points are i.i.d., we get

$$
\begin{equation*}
\mathcal{E}_{\rho}(\tau) \leq \mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}(A(\tau(s), Z))+\frac{2 L^{2}}{n} \operatorname{Tr}\left(\tau(s) \frac{X^{\top} X}{n}\right)\right], \tag{19}
\end{equation*}
$$

where $X \in \mathbb{R}^{n \times d}$ is the matrix with the inputs vectors $\left(x_{i}\right)_{i=1}^{n}$ as rows. The inequality above suggests us to introduce the surrogate problem

$$
\begin{equation*}
\min _{\tau \in \mathcal{T}} \hat{\mathcal{E}}_{\rho}(\tau) \quad \hat{\mathcal{E}}_{\rho}(\tau)=\mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}(A(\tau(s), Z))+\frac{2 L^{2}}{n} \operatorname{Tr}\left(\tau(s) \frac{X^{\top} X}{n}\right)\right], \tag{20}
\end{equation*}
$$

where, from the last inequality above, for any $\tau$, we have

$$
\begin{equation*}
\mathcal{E}_{\rho}(\tau) \leq \hat{\mathcal{E}}_{\rho}(\tau) \tag{21}
\end{equation*}
$$

We stress that the surrogate problem we take here is different from the one considered in previous work Bullins et al. (2019); Denevi et al. (2019a,b), where the authors considered as meta-objective only a part of the function above, namely, $\mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}(\mathcal{A}(\tau(s), Z))\right]$. As we will see in the following, such a choice is more appropriate for the problem at hand, since, differently from the meta-objective used in previous literature, it will allow us to develop a conditional meta-learning method that is theoretically grounded also for linear representation learning.

Exploiting Asm. 2, the surrogate problem in Eq. (20) can be restricted to the class of linear functions $\mathcal{T}_{\Phi}$ in Eq. (18) and it can be rewritten more explicitly as

$$
\begin{gather*}
\min _{H \in \mathcal{S}^{\mathrm{d} k}, \mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}}} \mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{\mathrm{Z} \sim \mu^{\mathrm{n}}} \mathcal{L}(\mathrm{H}, \mathrm{C}, \mathrm{~s}, \mathrm{Z}) \\
\mathcal{L}(\mathrm{H}, \mathrm{C}, \mathrm{~s}, \mathrm{Z})=  \tag{22}\\
\mathcal{R}_{\mathrm{Z}, \tau_{H}, \mathrm{C}(\mathrm{~s})}\left(\mathcal{A}\left(\tau_{H, \mathrm{C}}(\mathrm{~s}), \mathrm{Z}\right)\right)+\frac{2 \mathrm{~L}^{2}}{n} \operatorname{Tr}\left(\tau_{H, \mathrm{C}}(\mathrm{~s}) \frac{X^{\top} X}{n}\right) .
\end{gather*}
$$

In the following proposition we outline some useful properties of the meta-loss $\mathcal{L}(\cdot, \cdot, s, Z)$ introduced above (such as convexity) supporting its choice as surrogate meta-loss.

```
Algorithm 1 Meta-Algorithm, SGD on Eq. (22)
    Input \(\gamma>0\) meta-step size, \(\mathrm{H}_{0} \in \mathbb{S}_{+}^{d k}, \mathrm{C}_{0} \in \mathbb{S}_{+}^{d}\)
    Initialization \(H_{1}=H_{0} \in \mathbb{S}_{+}^{\mathrm{dk}}, \mathrm{C}=\mathrm{C}_{0} \in \mathbb{S}_{+}^{\mathrm{d}}\)
    For \(t=1\) to \(T\)
        Receive \(\left(\mu_{t}, s_{t}\right) \sim \rho\) and \(Z_{t} \sim \mu_{t}^{n}\)
        Let \(\theta_{\mathrm{t}}=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi\left(\mathrm{s}_{\mathrm{t}}\right)\right) \mathrm{H}_{\mathrm{t}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi\left(\mathrm{s}_{\mathrm{t}}\right)^{\top}\right)+\mathrm{C}_{\mathrm{t}}\)
        Compute \(\boldsymbol{w}_{\theta_{t}}=A\left(\theta_{t}, Z_{t}\right)\) by Eq. (2)
        Compute \(\nabla \mathcal{L}\left(\cdot, \mathrm{C}_{\mathrm{t}}, \mathrm{s}_{\mathrm{t}}, \mathrm{Z}_{\mathrm{t}}\right)\left(\mathrm{H}_{\mathrm{t}}\right)\) as in Eq. (23) with \(w_{\theta_{\mathrm{t}}}\)
        Compute \(\nabla \mathcal{L}\left(H_{t}, \cdot, s_{t}, Z_{t}\right)\left(\mathrm{C}_{\mathrm{t}}\right)\) as in Eq. (23) with \(w_{\theta_{t}}\)
        Update \(\mathrm{H}_{\mathrm{t}+1}=\operatorname{proj}_{\Theta}\left(\mathrm{H}_{\mathrm{t}}-\gamma \nabla \mathcal{L}\left(\cdot, \mathrm{C}_{\mathrm{t}}, \mathrm{s}_{\mathrm{t}}, \mathrm{Z}_{\mathrm{t}}\right)\left(\mathrm{H}_{\mathrm{t}}\right)\right)\)
        Update \(\mathrm{C}_{\mathrm{t}+1}=\operatorname{proj}_{\Theta}\left(\mathrm{C}_{\mathrm{t}}-\gamma \nabla \mathcal{L}\left(\mathrm{H}_{\mathrm{t}}, \cdot, \mathrm{s}_{\mathrm{t}}, \mathrm{Z}_{\mathrm{t}}\right)\left(\mathrm{C}_{\mathrm{t}}\right)\right)\)
    Return \(\overline{\mathrm{H}}=\frac{1}{\mathrm{~T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \mathrm{H}_{\mathrm{t}}, \overline{\mathrm{C}}=\frac{1}{\mathrm{~T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \mathrm{C}_{\mathrm{t}}\)
```

Proposition 4 (Properties of the Surrogate Meta-Loss $\mathcal{L}$ ). For any $\mathrm{Z} \in \mathcal{D}$ and $s \in \mathcal{S}$, the function $\mathcal{L}(\cdot, \cdot, s, Z)$ is convex and one of its subgradients is given, for any $\mathrm{H} \in \mathbb{S}_{+}^{d k}$ and $\mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}}$, by

$$
\begin{align*}
& \nabla \mathcal{L}(\mathrm{H}, \cdot, \mathrm{~s}, \mathrm{Z})(\mathrm{C})=\hat{\nabla} \\
& \nabla \mathcal{L}(\cdot, \mathrm{C}, \mathrm{~s}, \mathrm{Z})(\mathrm{H})=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right) \hat{\nabla}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right) \tag{23}
\end{align*}
$$

where

$$
\hat{\nabla}=-\frac{\lambda}{2} \tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger} w_{\tau_{\mathrm{H}, \mathrm{C}}(s)} w_{\tau_{\mathrm{H}, \mathrm{C}}(s)}^{\top} \tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger}+\frac{2 \mathrm{~L}^{2} \mathrm{X}^{\top} X}{\mathrm{n}^{2}} .
$$

Moreover, under Asm. 1 and Asm. 3, we have

$$
\|\nabla \mathcal{L}(\cdot, \cdot, s, Z)(H, C)\|_{F} \leq\left(1+K^{2}\right)(L R)^{2}\left(\frac{1}{2}+\frac{2}{n}\right) .
$$

The proof of Prop. 4 is reported in App. D.2. It follows from combining results from Denevi et al. (2019b) with the composition of the linear parametrization of the functions $\tau_{\mathrm{H}, \mathrm{C}} \in \mathcal{T}_{\Phi}$.

The Conditional Meta-Learning Estimator. The meta-learning strategy we propose consists in applying Stochastic Gradient Descent (SGD) on the surrogate problem in Eq. (22). Such a meta-algorithm is implemented in Alg. 1: we assume to observe a sequence of i.i.d. pairs $\left(Z_{t}, s_{t}\right)_{t=1}^{\top}$ of training datasets and side information, and at each iteration we update the conditional parameters $\left(\mathrm{H}_{\mathrm{t}}, \mathrm{C}_{\mathrm{t}}\right)$ by performing a step of constant size $\gamma>0$ in the direction of $-\nabla \mathcal{L}\left(\cdot, \cdot, s_{t}, Z_{t}\right)\left(\mathrm{H}_{\mathrm{t}}, \mathrm{C}_{\mathrm{t}}\right)$ and a projection step on $\mathbb{S}_{+}^{\mathrm{dk}} \times \mathbb{S}_{+}^{d}$. Finally, we output the conditioning function $\tau_{\overline{\mathrm{H}}, \overline{\mathrm{C}}}$ parametrized by $(\overline{\mathrm{H}}, \overline{\mathrm{C}})$, the average across all the iterates $\left(H_{t}, C_{t}\right)_{t=1}^{\top}$. The theorem below analyzes the generalization properties of such a conditioning function.

Theorem 5 (Excess Risk Bound for the Conditioning Function Returned by Alg. 1). Let Asm. 1 and Asm. 3 hold. For any $s \sim \rho_{\mathcal{S}}$, recall the conditional covariance matrices $W(s)$ and $\mathrm{C}(\mathrm{s})$ introduced in Thm. 1. Let $\tau_{\mathrm{H}, \mathrm{C}}$ be a fixed function in $\mathcal{T}_{\Phi}$ such that $\operatorname{Ran}(\mathrm{W}(\mathrm{s})) \subseteq$ $\operatorname{Ran}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{s})\right)$ for any $\mathrm{s} \sim \rho_{\mathcal{S}}$. Let $\overline{\mathrm{H}}$ and $\overline{\mathrm{C}}$ be the outputs of Alg. 1 applied to a sequence $\left(Z_{t}, s_{t}\right)_{t=1}^{\top}$ of i.i.d. pairs sampled from $\rho$ with meta-step size

$$
\begin{equation*}
\gamma=\frac{\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}}{\left(1+\mathrm{K}^{2}\right)(\mathrm{LR})^{2}}\left(\frac{1}{2}+\frac{2}{\mathrm{n}}\right)^{-1} \frac{1}{\sqrt{\mathrm{~T}}} \tag{24}
\end{equation*}
$$

Then, in expectation w.r.t. the sampling of $\left(Z_{t}, s_{t}\right)_{t=1}^{\top}$,

$$
\begin{aligned}
\mathbb{E} \mathcal{E}_{\rho}\left(\tau_{\bar{H}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \leq & \frac{\mathbb{E}_{s \sim \rho_{\mathcal{S}}} \operatorname{Tr}\left(\tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger} \mathrm{W}(\mathrm{~s})\right)}{2}+\frac{2 \mathrm{~L}^{2} \mathbb{E}_{\mathrm{s} \sim \rho_{\mathcal{S}}} \operatorname{Tr}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s}) \mathrm{C}(\mathrm{~s})\right)}{n} \\
& +\left(\frac{1}{2}+\frac{2}{n}\right) \frac{\left(1+\mathrm{K}^{2}\right)(\mathrm{LR})^{2}\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}}{\sqrt{\mathrm{~T}}}
\end{aligned}
$$

Proof (Sketch). The detailed proof is reported in App. D.4. Exploiting the fact that, for any $\tau \in \mathcal{T}, \mathcal{E}_{\rho}(\tau) \leq \widehat{\mathcal{E}}_{\rho}(\tau)$ (see Eq. (21)) and adding $\pm \widehat{\mathcal{E}}_{\rho}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)$, we can write the following

$$
\begin{equation*}
\mathbb{E}_{\mathbf{Z}} \mathcal{E}_{\rho}\left(\tau_{\overline{\mathrm{H}}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \leq \underbrace{\mathbb{E}_{\mathrm{Z}} \hat{\mathcal{E}}_{\rho}\left(\tau_{\overline{\mathrm{H}}, \overline{\mathrm{C}}}\right)-\hat{\mathcal{E}}_{\rho}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)}_{\mathrm{A}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)}+\underbrace{\hat{\mathcal{E}}_{\rho}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)-\mathcal{E}_{\rho}^{*}}_{\mathrm{B}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)} \tag{25}
\end{equation*}
$$

The term $\mathrm{A}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)$ can be controlled according to the convergence properties of the metaalgorithm in Alg. 1 as described in Prop. 12. Regarding the term $B\left(\tau_{H, C}\right)$, exploiting the definition of the within-task algorithm in Eq. (2) as minimum, for any $\tau \in \mathcal{T}$ such that $\operatorname{Ran}\left(\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} \mathcal{w}_{\mu} \mathcal{w}_{\mu}^{\top}\right) \subseteq \operatorname{Ran}(\tau(s))$ for any $s \sim \rho_{\mathcal{S}}$, we can rewrite

$$
\begin{aligned}
\mathrm{B}(\tau) & =\mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}(A(\tau(s), Z))-\mathcal{R}_{\mu}\left(w_{\mu}\right)\right]+\frac{2 \mathrm{~L}^{2} \mathbb{E}_{(\mu, s) \sim \rho} \operatorname{Tr}\left(\tau(s) \mathbb{E}_{x \sim \eta_{\mu}} x x^{\top}\right)}{n} \\
& \leq \frac{\mathbb{E}_{(\mu, s) \sim \rho} \operatorname{Tr}\left(\tau(s)^{\dagger} w_{\mu} w_{\mu}^{\top}\right)}{2}+\frac{2 \mathrm{~L}^{2} \mathbb{E}_{(\mu, s) \sim \rho} \operatorname{Tr}\left(\tau(s) \mathbb{E}_{x \sim \eta_{\mu}} x x^{\top}\right)}{n} .
\end{aligned}
$$

The desired statement then derives from combining the two parts above and optimizing w.r.t. $\gamma$.

We now present some important implications of Thm. 5.
Proposed Vs. Optimal Conditioning Function. Specializing the bound in Thm. 5 to the best conditioning function $\tau_{\rho}$ in Prop. 2, thanks to Asm. 2, we get the following bound for our estimator,

$$
\mathbb{E} \mathcal{E}_{\rho}\left(\tau_{\overline{\mathrm{H}}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \leq \mathcal{O}\left(\mathbb{E}_{\mathrm{s} \sim \rho_{\mathcal{S}}}\left\|\mathrm{W}(\mathrm{~s})^{1 / 2} \mathrm{C}(\mathrm{~s})^{1 / 2}\right\|_{*} \mathrm{n}^{-1 / 2}\right)+\mathcal{O}\left(\left\|\left(\mathrm{H}_{\rho}-\mathrm{H}_{0}, \mathrm{C}_{\rho}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}} \mathrm{~T}^{-1 / 2}\right)
$$

From such a bound, we can state that our proposed meta-algorithm achieves comparable performance to the best conditioning function $\tau_{\rho}$ in hindsight, when the number of observed tasks is sufficiently large. Moreover, recalling the unconditional oracle $\hat{\theta}_{\rho}$ in Eq. (15) used in previous literature, regarding the second term vanishing with $T$, we observe that our
conditional meta-learning approach incurs a cost of $\left\|\left(\mathrm{H}_{\rho}-\mathrm{H}_{0}, \mathrm{C}_{\rho}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}} \mathrm{T}^{-1 / 2}$ as opposed to the cost of $\left\|\hat{\theta}_{\rho}-\theta_{0}\right\| \mathrm{T}^{-1 / 4}$ associated to state-of-the-art unconditional meta-learning approaches (see Balcan et al. (2019); Bullins et al. (2019); Denevi et al. (2019b); Khodak et al. (2019)). Thus, our conditional approach presents a faster convergence rate w.r.t. T than such unconditional methods, but a complexity term that is expected to be larger due to the larger complexity of the class of functions we are working with. Such a faster rate w.r.t. T is essentially due to our formulation of the problem on the entire set of positivesemidefinite matrices (with no trace constraints). This in fact allows us to incorporate the within-task regularization parameter $\lambda$ directly in the linear representation and to gain a $\sqrt{T}$ order that was lost in previous literature when tuning w.r.t. the parameter $\lambda$. At the same time, this allows us to develop also a method requiring to tune just one hyper-parameter, while previous unconditional approaches requires to tune two hyper-parameters.

Comparison to Unconditional Meta-Learning. Specializing Thm. 5 to the best unconditional estimator $\tau_{\mathrm{H}, \mathrm{C}} \equiv \theta_{\rho}$ we introduced in Eq. (13), the bound for our estimator becomes

$$
\begin{equation*}
\mathbb{E} \mathcal{E}_{\rho}\left(\tau_{\overline{\mathrm{H}}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \leq \mathcal{O}\left(\left\|W_{\rho}^{1 / 2} \mathrm{C}_{\rho}^{1 / 2}\right\|_{*} n^{-1 / 2}\right)+\mathcal{O}\left(\left\|\theta_{\rho}-\mathrm{C}_{0}\right\| \mathrm{T}^{-1 / 2}\right) . \tag{26}
\end{equation*}
$$

From the bound above, we can conclude that the conditional approach provides, at least, the same guarantees as its unconditional counterpart. Moreover, we stress again that the bound above presents a faster rate w.r.t. T in comparison to the state-of-the-art unconditional methods.

Remark 3 (Online Variant of Eq. (2)). Also in this case, as already observed for the bias regularization and fine tuning framework proposed in Denevi et al. (2020), when we use the online inner family in Rem. 2, we can approximate the meta-subgradient in Eq. (23) by replacing the batch regularized empirical risk minimizer $\mathcal{A}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{s}), \mathrm{Z}\right)$ in Eq. (2) with the last iterate of the online algorithm in Eq. (4).

## 5 Experiments

We now present preliminary experiments in which we compare the proposed conditional meta-learning approach in Alg. 1 (cond.) with the unconditional counterpart (uncond.) and solving the tasks independently (ITL, namely, running the inner algorithm separately across the tasks with the constant linear representation $\theta=I_{d} \in \mathbb{S}_{+}^{d}$ ). We considered regression problems and we evaluated the errors by $\ell$ the absolute loss. We implemented the online variant of the within-task algorithm introduced in Eq. (4). The hyper-parameter $\gamma$ was chosen by (meta-)cross validation on separate $\mathrm{T}_{\mathrm{tr}}$, $\mathrm{T}_{\mathrm{va}}$ and $\mathrm{T}_{\mathrm{te}}$ respectively meta-train, -validation and -test sets. Each task is provided with a training dataset $Z_{\text {tr }}$ of $n_{\text {tr }}$ points and a test dataset $Z_{\text {te }}$ of $n_{\text {te }}$ points used to evaluate the performance of the within-tasks algorithm. In App. E we report the details of this process in our experiments.

Synthetic Clusters. We considered two variants of the setting described in Ex. 1 with side information corresponding to the training datasets $\mathrm{Z}_{\text {tr }}$ associated to each task. In


Figure 1: Test error (averaged over 5 random generations of the data) of different methods w.r.t. an increasing number of tasks on synthetic data. 2 clusters (Left) and 6 clusters (Right).
both settings, we sampled $T_{\text {tot }}=900$ tasks from a uniform mixture of $m$ clusters. For each task $\mu$, we generated the target vector $w_{\mu} \in \mathbb{R}^{\mathrm{d}}$ with $\mathrm{d}=20$ as $w_{\mu}=\mathrm{P}\left(\mathrm{j}_{\mu}\right) \tilde{w}_{\mu}$, where, $j_{\mu} \in\{1, \ldots, m\}$ denotes the cluster from which the task $\mu$ was sampled and with the components of $\tilde{w}_{\mu} \in \mathbb{R}^{d /(10)}$ sampled from the Gaussian distribution $\mathcal{G}(0,1)$ and then $\tilde{w}_{\mu}$ normalized to have unit norm, with $\mathrm{P}\left(\mathrm{j}_{\mu}\right) \in \mathbb{R}^{\mathrm{d} \times \mathrm{d} /(10)}$ a matrix with orthonormal columns. We then generated the corresponding dataset $\left(x_{i}, y_{i}\right)_{i=1}^{n_{\text {not }}}$ with $n_{\text {tot }}=80$ according to the linear equation $y=\left\langle x, w_{\mu}\right\rangle+\epsilon$, with $x$ sampled uniformly on the unit sphere in $\mathbb{R}^{d}$ and $\epsilon$ sampled from a Gaussian distribution, $\epsilon \sim \mathcal{G}(0,0.1)$. In this setting, the operator norm of the inputs' covariance matrix is small (equal to $1 / \mathrm{d}$ ) and the weight vectors' covariance matrix of each single cluster is low-rank (its rank is $d /(10)=2$ ). We implemented our conditional method using the feature map $\Phi: \mathcal{D} \rightarrow \mathbb{R}^{2 \mathrm{~d}}$ defined by $\Phi(Z)=\frac{1}{n_{\mathrm{tr}}} \sum_{i=1}^{n_{\mathrm{tr}}} \phi\left(z_{i}\right)$, with $\phi\left(z_{i}\right)=\operatorname{vec}\left(x_{i}\left(y_{i}, 1\right)^{\top}\right)$, where, for any matrix $A=\left[a_{1}, a_{2}\right] \in \mathbb{R}^{d \times 2}$ with columns $a_{1}, a_{2} \in \mathbb{R}^{d}, \operatorname{vec}(A)=\left(a_{1}, a_{2}\right)^{\top} \in \mathbb{R}^{2 d}$.

In Fig. 1, we report the results we got on an environment of tasks generated as above with $\mathfrak{m}=2$ (Left) and $\mathfrak{m}=6$ (Right) clusters, respectively. As we can see, when the clusters are two, the unconditional approach outperforms ITL (as predicted from previous literature), but the unconditional method is in turn outperformed by our conditional counterpart. When the number of clusters raises to six, the performance of unconditional meta-learning degrades to the same performance of ITL, while conditional meta-learning outperforms both methods. Summarizing, the more the heterogeneity of the environment (number of clusters) is significant, the more the conditional approach brings advantage w.r.t. the unconditional one. This is in line with our statement in Ex. 1.

Real Datasets. We tested the performance of the methods also on the regression problem on the computer survey data from Lenk et al. (1996) (see also McDonald et al., 2016). $T_{\text {tot }}=180$ people (tasks) rated the likelihood of purchasing one of $n_{\text {tot }}=20$ computers. The input represents $d=13$ computers' characteristics and the label is a rate in $\{0, \ldots, 10\}$. In this case, we used as side information the training datapoints $Z=\left(z_{i}\right)_{i=1}^{n_{r}}$ and the feature map $\Phi: \mathcal{D} \rightarrow \mathbb{R}^{\mathrm{d}+1}$ defined by $\Phi(Z)=w_{Z}$, with $w_{Z}$ the solution of Tikhonov regularization with the squared loss, namely, the vector satisfying $\left(\hat{X}^{\top} \widehat{X}+I_{d+1}\right) w_{Z}=\hat{X}^{\top} y$, where, $\widehat{X} \in \mathbb{R}^{(d+1) \times n}$ is the matrix obtained by adding to the matrix $X \in \mathbb{R}^{n \times d}$ one column of ones at the end. Fig. 2 (Left) shows that also in this case, the unconditional approach outperforms ITL, but the performance of its conditional counterpart is much better.


Figure 2: Test error (averaged over 5 random splitting of the data) of different methods w.r.t. an increasing number of tasks on the Lenk dataset (Left) and the Movielens-100k dataset (Right).


Figure 3: Test error (averaged over 5 random splitting of the data) of different methods w.r.t. an increasing number of tasks on the Jester-1 dataset.

Finally, we tested the performance of the methods on the Movielens-100k and Jester1 real-world datasets, containing ratings of users (tasks) to movies and jokes (points), respectively. We recall that recommendation system settings with ditems can be interpreted within the meta-learning setting by considering each data point $(x, y)$ to have input $x \in \mathbb{R}^{d}$ to be the one-hot encoding of the current item to be rated (e.g. a movie or a joke) and $y \in \mathbb{R}$ the corresponding score (see e.g. Denevi et al., 2019a, for more details). We restricted the original dataset to the $n_{\text {tot }}=20$ most voted movies/jokes (as a consequence, by formulation, $\mathrm{d}=20$ ). We guaranteed each user voted at least 5 movies/jokes, which led to a total of $\mathrm{T}_{\text {tot }}=400 / 450$ tasks (i.e. users). In both cases, we used as side information the training datapoints $Z=\left(z_{i}\right)_{i=1}^{n_{\text {rr }}}$. For the Movielens-100k dataset we used the same feature map described for the synthetic clusters experiments in Fig. 1. For the Jester-1 dataset, let $M$ and $m$ denote the maximum and minimum rating value that can be assigned to a joke. We adopted the feature map $\Phi: \mathcal{D} \rightarrow \mathbb{R}^{2 d+1}$ such that, for any dataset $Z=\left(x_{i}, y_{i}\right)_{i=1}^{n}$, we have

$$
\begin{equation*}
\Phi(Z)=\binom{\operatorname{vec}(\tilde{\Phi}(Z))}{1} \tag{27}
\end{equation*}
$$

where vec denotes the vectorization operator (i.e. mapping a matrix in the vector concate-
nating all its columns) and $\widetilde{\Phi}: Z \rightarrow \mathbb{R}^{\mathrm{d} \times 2}$ is such that

$$
\begin{equation*}
\widetilde{\Phi}(Z)=\left[\cos \left(\sum_{i=1}^{n} x_{i}\left(\frac{\pi}{4} \frac{M-y_{i}}{M-m}\right)\right), \sin \left(\sum_{i=1}^{n} x_{i}\left(\frac{\pi}{4} \frac{M-y_{i}}{M-m}\right)\right)\right] \odot\left(\sum_{i=1}^{n} x_{i}\right), \tag{28}
\end{equation*}
$$

with $\odot$ denoting the Hadamard (entry-wise) product broad-casted across both columns.
The rationale behind this feature map is to represent as similar vectors those users with similar scores for the same movies. In particular, each item-score pair observed in training is represented as a unitary vector in $\mathbb{R}_{++}^{2}$, with the angle depending on the score attributed to that item (the vector corresponds to zero if that movie was not observed at the training time). We noticed that this feature map did not provide significant advantages on the Movielens-100k dataset, while being particularly favorable on the Jester-1 benchmark.

We report the average test errors (and standard deviation) for ITL, conditional and unconditional meta-learning in Fig. 2 (Right) and Fig. 3 for Movielens-100k and Jester-1, respectively. As it can be noticed, the proposed approach performs significantly better than ITL and its unconditional counterpart. This suggests that groups of users might rely each on similar features (but different from those of other groups) to rate an item in the dataset (respectively a movie or a joke).

## 6 Conclusion

We proposed a conditional meta-learning approach aiming at learning a function mapping task's side information into a linear representation that is well suited for the task at hand. We theoretically and experimentally showed that the proposed conditional approach is advantageous w.r.t. the standard unconditional counterpart when the observed tasks share heterogeneous linear representations. Our investigation allowed us to develop also a new variant of an unconditional meta-learning method requiring tuning one less hyper-parameter and relying on faster learning bounds than state-of-the-art unconditional approaches.

We identify two main directions for future work. A first question left opened by most conditional meta-learning methods is how to design a suitable feature map $\Phi$ when the tasks' training datas is used as side information. Following most previous work Rusu et al. (2018); Wang et al. (2020) in our experiments we adopted a mean embedding representation. However, given the key importance played by such feature map in Thm. 5, it will be worth investigating better alternatives in the future. A second direction is more focused on computations and modeling aspects. In particular it will be valuable to investigate how to predict non-linear conditioning functions (similarly to e.g. Bertinetto et al. (2018); Finn et al. (2017)) and develop more efficient versions of our method, using less expensive algorithms to update the positive matrices, such as the Frank-Wolfe algorithm used in Bullins et al. (2019) to deal with unconditional settings.

## References

Abernethy, J., Bach, F., Evgeniou, T., and Vert, J.-P. (2009). A new approach to collaborative filtering: Operator estimation with spectral regularization. Journal of Machine Learning

Research, 10(Mar):803-826.
Argyriou, A., Evgeniou, T., and Pontil, M. (2008). Convex multi-task feature learning. Machine Learning, 73(3):243-272.

Balcan, M.-F., Khodak, M., and Talwalkar, A. (2019). Provable guarantees for gradientbased meta-learning. In International Conference on Machine Learning, pages 424-433.

Baxter, J. (2000). A model of inductive bias learning. J. Artif. Intell. Res., 12(149-198):3.
Bertinetto, L., Henriques, J. F., Torr, P. H., and Vedaldi, A. (2018). Meta-learning with differentiable closed-form solvers. arXiv preprint arXiv:1805.08136.

Bousquet, O. and Elisseeff, A. (2002). Stability and generalization. Journal of machine learning research, 2(Mar):499-526.

Bullins, B., Hazan, E., Kalai, A., and Livni, R. (2019). Generalize across tasks: Efficient algorithms for linear representation learning. In Algorithmic Learning Theory, pages 235-246.

Cai, T. T., Liang, T., and Rakhlin, A. (2020). Weighted message passing and minimum energy flow for heterogeneous stochastic block models with side information. Journal of Machine Learning Research, 21(11):1-34.

Caruana, R. (1997). Multitask learning. Machine learning, 28(1):41-75.
Denevi, G., Ciliberto, C., Grazzi, R., and Pontil, M. (2019a). Learning-to-learn stochastic gradient descent with biased regularization. In International Conference on Machine Learning, pages 1566-1575.

Denevi, G., Ciliberto, C., Stamos, D., and Pontil, M. (2018). Incremental learning-to-learn with statistical guarantees. In Proc. 34th Conference on Uncertainty in Artificial Intelligence (UAI).

Denevi, G., Pontil, M., and Ciliberto, C. (2020). The advantage of conditional meta-learning for biased regularization and fine tuning. Advances in Neural Information Processing Systems, 33.

Denevi, G., Stamos, D., Ciliberto, C., and Pontil, M. (2019b). Online-within-online metalearning. In Advances in Neural Information Processing Systems, pages 13089-13099.

Finn, C., Abbeel, P., and Levine, S. (2017). 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.

Finn, C. and Levine, S. (2018). Meta-learning and universality: Deep representations and gradient descent can approximate any learning algorithm. In International Conference on Learning Representations.

Finn, C., Rajeswaran, A., Kakade, S., and Levine, S. (2019). Online meta-learning. In International Conference on Machine Learning, pages 1920-1930.

Harper, F. M. and Konstan, J. A. (2015). The movielens datasets: History and context. Acm transactions on interactive intelligent systems (tiis), 5(4):1-19.

Hogben, L. (2006). Handbook of linear algebra. CRC press.
Hogben, L. (2013). Handbook of linear algebra. CRC Press.
Jacob, L., Vert, J.-p., and Bach, F. R. (2009). Clustered multi-task learning: A convex formulation. In Advances in neural information processing systems, pages 745-752.

Jerfel, G., Grant, E., Griffiths, T., and Heller, K. A. (2019). Reconciling meta-learning and continual learning with online mixtures of tasks. In Advances in Neural Information Processing Systems, pages 9119-9130.

Khodak, M., Balcan, M.-F. F., and Talwalkar, A. S. (2019). Adaptive gradient-based meta-learning methods. In Advances in Neural Information Processing Systems, pages 5915-5926.

Lenk, P. J., DeSarbo, W. S., Green, P. E., and Young, M. R. (1996). Hierarchical bayes conjoint analysis: Recovery of partworth heterogeneity from reduced experimental designs. Marketing Science, 15(2):173-191.

Maurer, A. (2009). Transfer bounds for linear feature learning. Machine Learning, 75(3):327-350.

Maurer, A., Pontil, M., and Romera-Paredes, B. (2013). Sparse coding for multitask and transfer learning. In International Conference on Machine Learning.

Maurer, A., Pontil, M., and Romera-Paredes, B. (2016). The benefit of multitask representation learning. The Journal of Machine Learning Research, 17(1):2853-2884.

McDonald, A. M., Pontil, M., and Stamos, D. (2016). New perspectives on k-support and cluster norms. Journal of Machine Learning Research, 17(155):1-38.

Micchelli, C. A., Morales, J. M., and Pontil, M. (2013). Regularizers for structured sparsity. Advances in Computational Mathematics, 38(3):455-489.

Pentina, A. and Lampert, C. (2014). A PAC-Bayesian bound for lifelong learning. In International Conference on Machine Learning, pages 991-999.

Rusu, A. A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., and Hadsell, R. (2018). Meta-learning with latent embedding optimization. arXiv preprint arXiv:1807.05960.

Shalev-Shwartz, S. and Ben-David, S. (2014). Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press.

Steinwart, I. and Christmann, A. (2008). Support vector machines. Springer Science \& Business Media.

Tripuraneni, N., Jin, C., and Jordan, M. I. (2020). Provable meta-learning of linear representations. arXiv preprint arXiv:2002.11684.

Vuorio, R., Sun, S.-H., Hu, H., and Lim, J. J. (2019). Multimodal model-agnostic metalearning via task-aware modulation. In Advances in Neural Information Processing Systems, pages 1-12.

Wang, R., Demiris, Y., and Ciliberto, C. (2020). A structured prediction approach for conditional meta-learning. Advances in Neural Information Processing Systems.

Yao, H., Wei, Y., Huang, J., and Li, Z. (2019). Hierarchically structured meta-learning. arXiv preprint arXiv:1905.05301.

## Appendix

The supplementary material is organized as follows. In App. A we give the bound on the generalization error of the algorithm in Eq. (2) that we used in various proofs. In App. B we report the proof to get the closed form of the best conditioning function $\tau_{\rho}$ outlined in Prop. 2. In App. C we report the proof of the statement in Ex. 1. 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. 5 for Alg. 1. Finally, in App. E we report the experimental details we missed in the main body.

## A Generalization Bound of the Within-Task Algorithm

We now study the generalization error of the within-task algorithm in Eq. (2), i.e. the discrepancy between the (true) risk and the empirical risk of the corresponding estimator. This is done in the following result where we exploit stability arguments, more precisely the so-called hypothesis stability, see (Bousquet and Elisseeff, 2002, Def. 3).

Proposition 6 (Generalization Error of the Within-Task Algorithm in Eq. (2)). Let Asm. 1 hold. For a distribution $\mu \sim \rho$, fix a dataset $Z=\left(x_{i}, y_{i}\right)_{i=1}^{n} \sim \mu^{n}$. For any $\theta \in \Theta$, let $w_{\theta}(Z)$ be the corresponding RERM in Eq. (2) over Z. Then, the following generalization error bound holds for $w_{\theta}(Z)$ :

$$
\begin{equation*}
\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{\mu}\left(w_{\theta}(Z)\right)-\mathcal{R}_{Z}\left(w_{\theta}(Z)\right)\right] \leq \frac{2 L^{2}}{n} \operatorname{Tr}\left(\mathbb{E}_{z \sim \mu} \theta x x^{\top}\right) . \tag{29}
\end{equation*}
$$

Proof. During this proof, we need to make explicit the dependency of the RERM (Regularized Empirical Risk Minimizer) $w_{\theta}$ in Eq. (2) w.r.t. the dataset Z. For any $i \in\{1, \ldots, n\}$, consider the dataset $Z^{(i)}$, a copy of the original dataset $Z$ in which we exchange the point $z_{i}=\left(x_{i}, y_{i}\right)$ with a new i.i.d. point $z_{i}^{\prime}=\left(x_{i}^{\prime}, y_{i}^{\prime}\right)$. For a fixed $\theta \in \Theta$, we analyze how much this perturbation affects the outputs of the RERM algorithm in Eq. (2). In other words, we study the discrepancy between $w_{\theta}(Z)$ and $w_{\theta}\left(Z^{(i)}\right)$. We start from observing that, since by Asm. $1 \mathcal{R}_{Z, \theta}$ is 1 -strongly convex w.r.t. $\|\cdot\|_{\theta}=\sqrt{\left\langle\cdot, \theta^{\dagger} \cdot\right\rangle}$, by growth condition and the definition of the RERM algorithm, we can write the following

$$
\begin{align*}
& \frac{1}{2}\left\|w_{\theta}\left(Z^{(i)}\right)-w_{\theta}(Z)\right\|_{\theta}^{2} \leq \mathcal{R}_{Z, \theta}\left(w_{\theta}\left(Z^{(i)}\right)\right)-\mathcal{R}_{Z, \theta}\left(w_{\theta}(Z)\right)  \tag{30}\\
& \frac{1}{2}\left\|w_{\theta}\left(Z^{(i)}\right)-w_{\theta}(Z)\right\|_{\theta}^{2} \leq \mathcal{R}_{Z^{(i)}, \theta}\left(w_{\theta}(Z)\right)-\mathcal{R}_{Z^{(i)}, \theta}\left(w_{\theta}\left(Z^{(i)}\right)\right) .
\end{align*}
$$

Hence, summing the two inequalities above, we get

$$
\begin{align*}
\left\|w_{\theta}\left(Z^{(i)}\right)-w_{\theta}(Z)\right\|_{\theta}^{2} & \leq \mathcal{R}_{Z, \theta}\left(w_{\theta}\left(Z^{(i)}\right)\right)-\mathcal{R}_{Z^{(i)}, \theta}\left(w_{\theta}\left(Z^{(i)}\right)\right)+\mathcal{R}_{Z^{(i)}, \theta}\left(w_{\theta}(Z)\right)-\mathcal{R}_{Z, \theta}\left(w_{\theta}(Z)\right) \\
& =\frac{A+B}{n}, \tag{31}
\end{align*}
$$

where we have introduced the terms

$$
\begin{align*}
& \mathrm{A}=\ell\left(\left\langle x_{i}^{\prime}, w_{\theta}(Z)\right\rangle, y_{i}^{\prime}\right)-\ell\left(\left\langle x_{i}^{\prime}, w_{\theta}\left(Z^{(i)}\right)\right\rangle, y_{i}^{\prime}\right)  \tag{32}\\
& \mathrm{B}=\ell\left(\left\langle x_{i}, w_{\theta}\left(Z^{(i)}\right)\right\rangle, y_{i}\right)-\ell\left(\left\langle x_{i}, w_{\theta}(Z)\right\rangle, y_{i}\right) .
\end{align*}
$$

Now, introducing the subgradients $s_{\theta, i}^{\prime} \in \partial \ell\left(\cdot, y_{i}^{\prime}\right)\left(\left\langle x_{i}^{\prime}, w_{\theta}(Z)\right\rangle\right)$ and $s_{\theta, i} \in \partial l\left(\cdot, y_{i}\right)\left(\left\langle x_{i}, w_{\theta}\left(Z^{(i)}\right)\right\rangle\right)$ and applying Holder's inequality, we can write

$$
\begin{align*}
& \mathrm{A} \leq\left\langle x_{i}^{\prime} s_{\theta, i}^{\prime}, w_{\theta}(Z)-w_{\theta}\left(Z^{(i)}\right)\right\rangle \leq\left\|x_{i}^{\prime} s_{\theta, i}^{\prime}\right\|_{\theta, *}\left\|w_{\theta}\left(Z^{(i)}\right)-w_{\theta}(Z)\right\|_{\theta}  \tag{33}\\
& \mathrm{B} \leq\left\langle x_{i} s_{\theta, i}, w_{\theta}\left(Z^{(i)}\right)-w_{\theta}(Z)\right\rangle \leq\left\|x_{i} s_{\theta, i}\right\|_{\theta, *}\left\|w_{\theta}\left(Z^{(i)}\right)-w_{\theta}(Z)\right\|_{\theta},
\end{align*}
$$

where $\|\cdot\|_{\theta, *}=\sqrt{\langle\cdot, \theta \cdot\rangle}$ is the dual norm of $\|\cdot\|_{\theta}$. Combining these last two inequalities with Eq. (31) and simplifying, we get the following

$$
\begin{equation*}
\left\|w_{\theta}\left(Z^{(i)}\right)-w_{\theta}(Z)\right\|_{\theta} \leq \frac{1}{n}\left(\left\|x_{i}^{\prime} s_{\theta, i}^{\prime}\right\|_{\theta, *}+\left\|x_{i} s_{\theta, i}\right\|_{\theta, *}\right) . \tag{34}
\end{equation*}
$$

Hence, combining the first row in Eq. (33) with Eq. (34), we can write

$$
\begin{equation*}
\ell\left(\left\langle x_{i}^{\prime}, w_{\theta}(Z)\right\rangle, y_{i}^{\prime}\right)-\ell\left(\left\langle x_{i}^{\prime}, w_{\theta}\left(Z^{(i)}\right)\right\rangle, y_{i}^{\prime}\right) \leq \frac{1}{n}\left(\left\|x_{i}^{\prime} s_{\theta, i}^{\prime}\right\|_{\theta, *}^{2}+\left\|x_{i}^{\prime} s_{\theta, i}^{\prime}\right\|_{\theta, *}\left\|x_{i} s_{\theta, i}\right\|_{\theta, *}\right) . \tag{35}
\end{equation*}
$$

Now, taking the expectation w.r.t. $Z \sim \mu^{n}$ and $z_{i}^{\prime} \sim \mu$ of the left side member above, according to (Bousquet and Elisseeff, 2002, Lemma 7), we get
$\mathbb{E}_{Z \sim \mu^{n}} \mathbb{E}_{z_{i}^{\prime} \sim \mu}\left[\ell\left(\left\langle x_{i}^{\prime}, w_{\theta}(Z)\right\rangle, y_{i}^{\prime}\right)-\ell\left(\left\langle x_{i}^{\prime}, w_{\theta}\left(Z^{(i)}\right)\right\rangle, y_{i}^{\prime}\right)\right]=\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{\mu}\left(w_{\theta}(Z)\right)-\mathcal{R}_{Z}\left(w_{\theta}(Z)\right)\right]$.
Finally, taking the expectation of the right side member, exploiting the fact that the points are i.i.d. according $\mu$, we get

$$
\begin{equation*}
\mathbb{E}_{Z \sim \mu^{n}} \mathbb{E}_{z_{i}^{\prime} \sim \mu} \frac{1}{n}\left(\left\|x_{i}^{\prime} s_{\theta, i}^{\prime}\right\|_{\theta, *}^{2}+\left\|x_{i}^{\prime} s_{\theta, i}^{\prime}\right\|_{\theta, *}\left\|x_{i} s_{\theta, i}\right\|_{\theta, *}\right) \leq \frac{2}{n} \mathbb{E}_{Z \sim \mu^{n}} \mathbb{E}_{z_{i}^{\prime} \sim \mu}\left\|x_{i}^{\prime} s_{\theta,}^{\prime}\right\|_{\theta, *}^{2} \tag{36}
\end{equation*}
$$

where we recall that $s_{\theta, i}^{\prime} \in \partial \ell\left(\cdot, y_{i}^{\prime}\right)\left(\left\langle x_{i}^{\prime}, w_{\theta}(Z)\right\rangle\right)$. Combining the two last statements above, we get

$$
\begin{equation*}
\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{\mu}\left(w_{\theta}(Z)\right)-\mathcal{R}_{Z}\left(w_{\theta}(Z)\right)\right] \leq \frac{2}{n} \mathbb{E}_{Z \sim \mu^{n}} \mathbb{E}_{z_{i}^{\prime} \sim \mu}\left\|x_{i}^{\prime} s_{\theta, i}^{\prime}\right\|_{\theta, *}^{2} \tag{37}
\end{equation*}
$$

Finally, substituting the close form of $\|\cdot\|_{\theta, *}$ and observing that, by Asm. 1 we have $\left\|x_{i}^{\prime} s_{\theta,}^{\prime}\right\|_{\theta, *}^{2} \leq \mathrm{L}^{2}\left\|x_{i}^{\prime}\right\|_{\theta, *}^{2}$, we get the desired statement:

$$
\begin{equation*}
\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{\mu}\left(w_{\theta}(Z)\right)-\mathcal{R}_{Z}\left(w_{\theta}(Z)\right)\right] \leq \frac{2 L^{2}}{n} \mathbb{E}_{z_{i}^{\prime} \sim \mu}\left\langle x_{i}^{\prime}, \theta x_{i}^{\prime}\right\rangle=\frac{2 L^{2}}{n} \operatorname{Tr}\left(\mathbb{E}_{z \sim \mu} \theta x x^{\top}\right) . \tag{38}
\end{equation*}
$$

## B Proof of Prop. 2

In this section we report the proof to get the closed form of the best conditioning function $\tau_{\rho}$ outlined in Prop. 2. In order to do this, we need the following results.

Lemma 7. For any $\mu \sim \rho_{\mathcal{M}}$, define the inputs' covariance matrix $\mathrm{C}_{\mu}=\mathbb{E}_{x \sim \eta_{\mu}} x \chi^{\top}$. Then, for any $w_{\mu} \in \operatorname{argmin}_{w \in \mathbb{R}^{d}} \mathcal{R}_{\mu}(w)$, the projection $w_{0, \mu}=\mathrm{C}_{\mu}^{\dagger} \mathrm{C}_{\mu} w_{\mu}$ of $w_{\mu}$ onto the range of $\mathrm{C}_{\mu}$ is still a minimizer of $\mathcal{R}_{\mu}$.

Proof. Consider the decomposition of $w_{\mu}$ w.r.t. the range of $\mathrm{C}_{\mu}$ :

$$
\begin{equation*}
w_{\mu}=w_{0, \mu}+w^{\perp} \tag{39}
\end{equation*}
$$

with $w_{0, \mu}=C_{\mu}^{\dagger} C_{\mu} w_{\mu}$ and $w^{\perp} \in \mathbb{R}^{d}$ such that $C_{\mu} w^{\perp}=0$. We note that, almost surely w.r.t. the points $x \in \mathbb{R}^{\mathrm{d}}$ sampled from $\mu$, we have $\left\langle w^{\perp}, x\right\rangle=0$. This follows by noting that by the orthogonality between $C_{\mu}$ and $w^{\perp}$, we have

$$
\begin{equation*}
0=\left\langle w^{\perp}, C_{\mu} w^{\perp}\right\rangle=\mathbb{E}_{x \sim \eta_{\mu}}\left\langle w^{\perp}, x x^{\top} w^{\perp}\right\rangle=\mathbb{E}_{x \sim \eta_{\mu}}\left\langle x, w^{\perp}\right\rangle^{2}, \tag{40}
\end{equation*}
$$

that can hold only if $\left\langle x, w^{\perp}\right\rangle^{2}=0$ almost surely (a.s.) w.r.t. $\eta_{\mu}$. We conclude that $\left\langle w_{\mu}, x\right\rangle=$ $\left\langle w_{0, \mu}, x\right\rangle+\left\langle w^{\perp}, x\right\rangle=\left\langle w_{0, \mu}, x\right\rangle$ a.s. w.r.t. $\mu$ and, consequently, $\mathcal{R}_{\mu}\left(w_{\mu}\right)=\mathcal{R}_{\mu}\left(w_{0, \mu}\right)$.

Corollary 8. For any $s \in \mathcal{S}$, recall the conditional covariance matrices in Thm. 1. Then, $\operatorname{Ran}(W(s)) \subset \operatorname{Ran}(C(s))$, namely the range of the task-vector conditional covariance $W(s)$ is always contained in the range of the input conditional covariance $\mathrm{C}(\mathrm{s})$.

Proof. The corollary is a direct consequence of the previous Lemma 7. The result above guarantees that for any $\mu \sim \rho_{\mathcal{M}}$, the rank-one operator $W_{\mu}=w_{\mu} \mathcal{w}_{\mu}^{\top}$ has range contained in the range of $C_{\mu}$. Taking the conditional expectations $W(s)=\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} W_{\mu}$ and $C(s)=$ $\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} C_{\mu}$ maintains this relation unaltered, giving the desired statement.

Lemma 9. Let $\mathrm{P} \in \mathbb{S}_{+}^{d}$ be an orthogonal projector, namely such that $\mathrm{P}=\mathrm{P}^{2}$. Then, for any positive definite matrix $\theta \in \mathbb{S}_{++}^{d}$, we have $\mathrm{P} \theta^{-1} \mathrm{P} \succeq(\mathrm{P} \theta \mathrm{P})^{\dagger}$.

Proof. The proof is essentially a corollary of Schur's complement. Let consider the decomposition

$$
\begin{equation*}
\theta=\underbrace{P \theta P}_{A}+\underbrace{P \theta(I-P)}_{B}+\underbrace{(I-P) \theta P}_{B^{\top}}+\underbrace{(I-P) \theta(I-P)}_{C} \tag{41}
\end{equation*}
$$

where $A, C \in \mathbb{S}_{+}^{d}, B \in \mathbb{R}^{d \times d}$ and $C B=B^{\top} C=0$ since $(I-P) P=P(I-P)=P-P^{2}=$ $\mathrm{P}-\mathrm{P}=0$. Additionally, since $\mathrm{C}^{\dagger}=C \mathrm{C}^{\dagger} \mathrm{C}^{\dagger}=\mathrm{C}^{\dagger} \mathrm{C}^{\dagger} \mathrm{C}$, we have that also $A C^{\dagger}=A C C^{\dagger} C^{\dagger}=0$ and analogously $\mathrm{C}^{\dagger} \mathrm{B}=\mathrm{B}^{\top} \mathrm{C}^{\dagger}=0$. Note that since $\theta$ is invertible, both A and C are full rank. We now observe a few relevant interactions between the objects above. In particular, we observe that $\mathrm{CC}^{\dagger} \mathrm{B}^{\top}=\mathrm{B}^{\top}$. To see this, first note that

$$
\begin{equation*}
C C^{\dagger} \mathrm{B}^{\top}=(\mathrm{I}-\mathrm{P}) \theta(\mathrm{I}-\mathrm{P})((\mathrm{I}-\mathrm{P}) \theta(\mathrm{I}-\mathrm{P}))^{\dagger}(\mathrm{I}-\mathrm{P}) \theta \mathrm{P} . \tag{42}
\end{equation*}
$$

By taking $\mathrm{D}=(\mathrm{I}-\mathrm{P}) \theta^{1 / 2}$ and using the properties of the pseudoinverse (e.g. $\mathrm{D}=$ $\left.\mathrm{D}^{\top}\left(\mathrm{DD}^{\top}\right)^{\dagger}\right)$, we have

$$
\begin{align*}
\mathrm{CC}^{\dagger} \mathrm{B}^{\top} & =\mathrm{DD}^{\top}\left(\mathrm{DD}^{\top}\right)^{\dagger} \mathrm{D} \theta^{1 / 2} \mathrm{P}  \tag{43}\\
& =\mathrm{DD}^{\dagger} D \theta^{1 / 2} \mathrm{P}  \tag{44}\\
& =\mathrm{D} \theta^{1 / 2} \mathrm{P}  \tag{45}\\
& =\mathrm{B}^{\top} . \tag{46}
\end{align*}
$$

We now derive an alternative characterization of $\theta$ in terms of $A, B, C$. By adding and removing a term $\mathrm{BC}^{\dagger} \mathrm{B}$ to $\theta$, we have

$$
\begin{align*}
\theta & =A+B+B^{\top}+C  \tag{47}\\
& =A-B C^{\dagger} B^{\top}+B C^{\dagger} B^{\top}+B+B^{\top}+C  \tag{48}\\
& =A-B C^{\dagger} B^{\top}+B+C+(B+C)\left(C^{\dagger} B^{\top}\right)  \tag{49}\\
& =A-B C^{\dagger} B^{\top}+B+C+\left(A-B C^{\dagger} B^{\top}+B+C\right)\left(C^{\dagger} B\right)  \tag{50}\\
& =\left(A-B C^{\dagger} B^{\top}+B+C\right)\left(I+C^{\dagger} B^{\top}\right), \tag{51}
\end{align*}
$$

where we have first used the equality $C C^{\dagger} \mathrm{B}^{\top}=\mathrm{B}^{\top}$ and then the ortogonality $A C^{\dagger}=$ $\mathrm{B}^{\top} \mathrm{C}^{\dagger}=0$. Following a similar reasoning

$$
\begin{align*}
A-B C^{\dagger} B^{\top}+B+C & =A-B C^{\dagger} B^{\top}+C+B C^{\dagger} C  \tag{52}\\
& =A-B C^{\dagger} B^{\top}+C+B C^{\dagger}\left(A-B C^{\dagger} B^{\top}+C\right)  \tag{53}\\
& =\left(I+B C^{\dagger}\right)\left(A-B C^{\dagger} B^{\top}+C\right) \tag{54}
\end{align*}
$$

since $\mathrm{BC}^{\dagger} \mathrm{C}=\mathrm{C}$ (following the same reasoning used for $\mathrm{B}^{\top}=\mathrm{CC}^{\dagger} \mathrm{B}^{\top}$ ) and $\mathrm{AC} C^{\dagger}=\mathrm{C}^{\dagger} \mathrm{B}=0$. We conclude that

$$
\begin{equation*}
\theta=\left(I+B C^{\dagger}\right)\left(A-B C^{\dagger} B^{\top}+C\right)\left(I+C^{\dagger} B^{\top}\right) . \tag{55}
\end{equation*}
$$

We now show that all terms in the equation above are invertible. First note that ( $\mathrm{I}+$ $\left.B C^{\dagger}\right)^{-1}=\left(I-B C^{\dagger}\right)$ and $\left(I+C^{\dagger} B^{\top}\right)^{-1}=\left(I+C^{\dagger} B^{\top}\right)$. Moreover, since $\theta \succ 0$ and $C(A-$ $\left.B C^{\dagger} B^{\top}\right)=0$, then also $A-B C^{\dagger} B^{\top} \succ 0$. We have

$$
\begin{equation*}
\theta^{-1}=\left(I-C^{\dagger} B^{\top}\right)\left(A-B C^{\dagger} B^{\top}+C\right)^{-1}\left(I-B C^{\dagger}\right) \tag{56}
\end{equation*}
$$

from which we conclude

$$
\begin{align*}
P \theta^{-1} \mathrm{P} & =\mathrm{P}\left(A-B C^{\dagger} \mathrm{B}^{\top}+C\right)^{-1} P  \tag{57}\\
& =P\left(\left(A-B C^{\dagger} B^{\top}\right)^{\dagger}+C^{\dagger}\right) P  \tag{58}\\
& =P\left(A-B C^{\dagger} B^{\top}\right)^{\dagger} P  \tag{59}\\
& =\left(A-B C^{\dagger} B^{\top}\right)^{\dagger} . \tag{60}
\end{align*}
$$

Since $B C^{\dagger} B^{\top} \succeq 0$, we have $A-B C^{\dagger} B^{\top} \preceq A$ and therefore $\left(A-B C^{\dagger} B^{\top}\right)^{\dagger} \succeq A^{\dagger}$ from which we have

$$
\begin{equation*}
P \theta^{-1} P=\left(A-B C^{\dagger} B^{\top}\right)^{\dagger} \succeq A^{\dagger}=(P \theta P)^{\dagger} \tag{61}
\end{equation*}
$$

as desired.

Proposition 10. Consider two matrices $A, B \in \mathbb{S}_{+}^{d}$ such that $\operatorname{Ran}(A) \subseteq \operatorname{Ran}(B)$ and consider the following associated problem:

$$
\begin{equation*}
\left.\min _{\theta \in \mathbb{S}_{+}^{d},}, \operatorname{Ran}(A) \subseteq \operatorname{Ran}(\theta)\right] \operatorname{Tr}\left(\theta^{-1} A\right)+\operatorname{Tr}(\theta B) . \tag{62}
\end{equation*}
$$

Then, a minimizer and the corresponding minimum of the problem above are given by

$$
\begin{equation*}
\theta_{*}=B^{-1 / 2}\left(B^{1 / 2} A B^{1 / 2}\right)^{1 / 2} B^{-1 / 2} \quad 2\left\|B^{1 / 2} A^{1 / 2}\right\|_{*} . \tag{63}
\end{equation*}
$$

Moreover $\theta_{*}$ is the unique minimizer such that $\operatorname{Ran}\left(\theta_{*}\right) \subset \operatorname{Ran}(\mathrm{B})$.
Proof. Let $\Theta=\left\{\theta \in \mathbb{S}_{+}^{d} \mid \operatorname{Ran}(A) \subset \operatorname{Ran}(\theta)\right\}$ and denote by $F: \Theta \rightarrow \mathbb{R}$ the objective functional of the problem in Eq. (62), such that for any $\theta \in \Theta$

$$
\begin{equation*}
F(\theta)=\operatorname{Tr}\left(\theta^{-1} A\right)+\operatorname{Tr}(\theta B) . \tag{64}
\end{equation*}
$$

Note that the sign of inverse is well defined since $\operatorname{Ran}(\mathcal{A}) \subset \operatorname{Ran}(\theta)$. We begin the proof by showing that the Eq. (62) is equivalent to

$$
\begin{equation*}
\min _{\theta \in \mathbb{S}_{+}^{\mathrm{d}}, \operatorname{Ran}(A) \subset \operatorname{Ran}(\theta) \subset \operatorname{Ran}(B)} \operatorname{Tr}\left(\theta^{-1} A\right)+\operatorname{Tr}(\theta B) . \tag{65}
\end{equation*}
$$

To see this, let $\mathrm{P}=\mathrm{BB}^{\dagger}$ the orthogonal projector onto the range of B . By hypothesis, $A=P A P$ and $B=P B P$. Therefore, for any $\theta \in \mathbb{S}_{++}^{d}$

$$
\begin{aligned}
F(\theta) & =\operatorname{Tr}\left(\theta^{-1} A\right)+\operatorname{Tr}(\theta B) \\
& =\operatorname{Tr}\left(P \theta^{-1} \mathrm{PA}\right)+\operatorname{Tr}(\mathrm{P} \theta \mathrm{~PB}) \\
& \geq \operatorname{Tr}\left((\mathrm{P} \theta \mathrm{P})^{\dagger} A\right)+\operatorname{Tr}(\mathrm{P} \theta \mathrm{~PB}) \\
& =\mathrm{F}(\mathrm{P} \theta \mathrm{P}),
\end{aligned}
$$

where we have applied the fact that $\mathrm{P} \theta^{-1} \mathrm{P} \succeq(\mathrm{P} \theta \mathrm{P})^{\dagger}$ from Lemma 9 and the positive semidefinteness of $A$. The inequality above implies the equivalence between Eq. (62) and Eq. (65). Indeed, let $\theta_{*} \in \Theta$ be a minimizer of Eq. (62) and consider a sequence $\left(\theta_{n}\right)_{n \in \mathbb{N}}$ such that $\theta_{n} \in \mathbb{S}_{++}^{d}$ for any $n \in \mathbb{N}$ and $\theta_{n} \rightarrow \theta_{*}$. By continuity of $F$ we have also that $F\left(\theta_{n}\right) \rightarrow F\left(\theta_{*}\right)$. Clearly, $F\left(\theta_{*}\right) \leq F\left(P \theta_{n} P\right) \leq F\left(\theta_{n}\right)$ and therefore also $F\left(P \theta_{n} P\right) \rightarrow F\left(\theta_{*}\right)$. By continuity of $F$ over $\Theta$, this also implies that the limit $\lim _{n \rightarrow+\infty} P \theta_{n} P=P \theta_{*} P$ is a minimizer for Eq. (62) (and one such that $\operatorname{Ran}\left(\theta_{*}\right) \subset \operatorname{Ran}(\mathrm{B})$ ). We consider now the set $\Theta_{\mathrm{B}}=\left\{\theta \in \mathbb{S}_{+}^{\mathrm{d}} \mid \operatorname{Ran}(\theta)=\operatorname{Ran}(\mathrm{B})\right\}$ of all positive semidefinite matrices with same range as $B$, hence invertible on $\operatorname{Ran}(B)$. Note that $\Theta_{B}$ is an open subset of $\Theta$ and its closure in $\Theta$ corresponds to $\Theta$ itself. By definition, any $\theta \in \Theta_{B}$ is such that $\theta=B^{\dagger / 2} X B^{\dagger / 2}$ with $\operatorname{Ran}(X)=\operatorname{Ran}(B)$. This implies in particular that $X B^{\dagger} B=X$ and $\theta^{\dagger}=B^{\dagger / 2} X^{\dagger} B^{\dagger / 2}$. Therefore,

$$
\begin{align*}
F(\theta) & =\operatorname{Tr}\left(\theta^{\dagger} A\right)+\operatorname{Tr}(\theta B)  \tag{66}\\
& =\operatorname{Tr}\left(X^{\dagger} B^{1 / 2} A B^{1 / 2}\right)+\operatorname{Tr}(X), \tag{67}
\end{align*}
$$

and $\operatorname{Ran}\left(\mathrm{B}^{1 / 2} A \mathrm{~B}^{1 / 2}\right) \subseteq \operatorname{Ran}(\mathrm{B})=\operatorname{Ran}(\mathrm{X})$. We can now minimize the problem w.r.t. $X$, namely

$$
\begin{equation*}
\min _{X \in \mathbb{S}_{+}^{d}, \operatorname{Ran}\left(B^{1 / 2} A B^{1 / 2}\right) \subseteq \operatorname{Ran}(X)} \operatorname{Tr}\left(X^{\dagger} B^{1 / 2} A B^{1 / 2}\right)+\operatorname{Tr}(X) . \tag{68}
\end{equation*}
$$

The minimization corresponds to the variational form of the trace norm of $B^{1 / 2} A B^{1 / 2}$ Micchelli et al. (2013) and has solution $X_{*}=\left(B^{1 / 2} A B^{1 / 2}\right)^{1 / 2}$, with minimum corresponding to $2 \operatorname{Tr}\left(\left(B^{1 / 2} A B^{1 / 2}\right)^{1 / 2}\right)=2\left\|B^{1 / 2} A^{1 / 2}\right\|_{*}$. To conclude the proof, let $G:\{X \in$ $\left.\mathbb{S}_{+}^{\mathrm{d}} \mid \operatorname{Ran}\left(\mathrm{B}^{1 / 2} \mathrm{AB}{ }^{1 / 2}\right) \subseteq \operatorname{Ran}(\mathrm{X})\right\} \rightarrow \mathbb{R}$ be the objective functional in Eq. (68) such that $G(X)=\operatorname{Tr}\left(X^{\dagger} B^{1 / 2} A B^{1 / 2}\right)+\operatorname{Tr}(X)$. Let now $X_{*} \in \mathbb{S}_{+}^{d}$ be a minimzer for $G$ and $\left(X_{n}\right)_{n \in \mathbb{N}}$ be a minimizing sequence with $\operatorname{Ran}\left(X_{n}\right)=\operatorname{Ran}(B)$ for each $n \in \mathbb{N}$ and $X_{n} \rightarrow X_{*}$. Let $\left(\theta_{n}\right)_{n \in \mathbb{N}}$ such that $\theta_{n}=B^{\dagger / 2} X B^{\dagger / 2}$ for any $n \in \mathbb{N}$. Then we have $\theta_{n} \rightarrow B^{\dagger / 2} X_{*} B^{\dagger / 2}$ and by continuity $F\left(B^{\dagger / 2} X_{*} B^{\dagger / 2}\right)=G\left(X_{*}\right)$, hence $\min _{X} G(X) \leq \min _{\theta} F(\theta)$. Note that $B^{\dagger / 2} X_{*} B^{\dagger / 2}$ is a minimizer for $F$, since $F$ and $G$ have same minimum value. To see this it is sufficient to show that, given a minimizing sequence $\left(\theta_{n}\right)_{n \in \mathbb{N}}$ such that $\operatorname{Ran}\left(\theta_{n}\right)=\operatorname{Ran}(B)$ for any $n \in \mathbb{N}$ and $\theta_{n} \rightarrow \theta_{*}$, we have $X_{n}=B^{1 / 2} \theta_{n} B^{1 / 2} \rightarrow B^{1 / 2} \theta_{n} B^{1 / 2}$ and thus $F\left(\theta_{*}\right)=G\left(B^{1 / 2} \theta_{n} B^{1 / 2}\right)$. We have shown that $\min _{\theta} F(\theta) \geq \min _{X} G(X)$. Therefore $\theta_{*}=B^{\dagger / 2} X_{*} B^{\dagger / 2}=B^{\dagger / 2}\left(B^{1 / 2} A B^{1 / 2}\right)^{1 / 2} B^{\dagger / 2}$ is a minimizer of Eq. (62) as desired. The uniqueness of $\theta_{*}$ follows from the uniqueness of $X_{*}$ from the standard results on the variational form of the trace norm Micchelli et al. (2013).

We now have all the ingredients necessary to prove Prop. 2.
Proposition 2 (Best Conditioning Function in Hindsight). The conditioning function minimizer and the minimum of the bound presented in Thm. 1 over the set

$$
\left\{\tau \in \mathcal{T} \quad \mid \quad \operatorname{Ran}(W(s)) \subseteq \operatorname{Ran}(\tau(s)), \rho_{\mathcal{S}} \text {-almost surely }\right\}
$$

are respectively

$$
\tau_{\rho}(s)=\frac{\sqrt{n}}{2 L} C(s)^{\dagger / 2}\left(C(s)^{1 / 2} W(s) C(s)^{1 / 2}\right)^{1 / 2} C(s)^{\dagger / 2}
$$

and

$$
\begin{equation*}
\mathcal{E}_{\rho}\left(\tau_{\rho}\right)-\mathcal{E}_{\rho}^{*} \leq \frac{2 L \mathbb{E}_{s \sim \rho_{\mathcal{S}}}\left\|W(s)^{1 / 2} \mathrm{C}(s)^{1 / 2}\right\|_{*}}{\sqrt{\mathrm{n}}} \tag{11}
\end{equation*}
$$

Proof. We aim to minimize

$$
\begin{equation*}
\min _{\substack{\tau: \mathcal{S} \rightarrow \Theta \\ W(s)) \subseteq \operatorname{Ran}(\tau(s))}} \mathbb{E}_{s \sim \rho_{\mathcal{S}}} \varphi(s, \tau(s)) \quad \text { with } \quad \varphi(s, \theta)=\frac{\operatorname{Tr}\left(\theta^{\dagger} \mathrm{W}(s)\right)}{2}+\frac{2 \mathrm{~L}^{2} \operatorname{Tr}(\theta \mathrm{C}(s))}{n} \tag{69}
\end{equation*}
$$

over the set of all measurable functions $\tau: \mathcal{S} \rightarrow \Theta$. Note that from Cor. 8 , for any $s \in \mathcal{S}$ we have $\operatorname{Ran}(W(s)) \subset \operatorname{Ran}(C(s))$. Therefore we can apply Prop. 10 to have that for any $s \in \mathcal{S}$, the problem

$$
\begin{equation*}
\min _{\theta \in \mathbb{S}_{+}^{\mathrm{d}}, \operatorname{Ran}(W(s)) \subseteq \operatorname{Ran}(\theta)} \varphi(s, \theta) \tag{70}
\end{equation*}
$$

has solution

$$
\begin{equation*}
\tau_{\rho}(s)=\frac{\sqrt{n}}{2 L} C(s)^{\dagger / 2}\left(C(s)^{1 / 2} W(s) C^{1 / 2}\right)^{1 / 2} C(s)^{\dagger / 2} \tag{71}
\end{equation*}
$$

Therefore, for any $\tau: \mathcal{S} \rightarrow \Theta$ we have

$$
\begin{equation*}
\mathbb{E}_{s \sim \rho_{\mathcal{S}}} \varphi\left(\tau_{\rho}(s), s\right) \leq \mathbb{E}_{s \sim \rho_{\mathcal{S}}} \varphi(\tau(s), s), \tag{72}
\end{equation*}
$$

and therefore $\mathbb{E}_{s \sim \rho_{\mathcal{S}}} \varphi\left(\tau_{\rho}(s), s\right) \leq \min _{\tau} \mathbb{E}_{s \sim \rho_{\mathcal{S}}} \varphi(\tau)$. To conclude the proof we need to show that $\tau_{\rho}$ is measurable. This follows immediately by applying Aumann's measurable selection principle, see for instance the formulation in (Steinwart and Christmann, 2008, Lemma A.3.18). Under the notation of Steinwart and Christmann (2008), we can apply the result by taking $h(s, \theta)=\left(\theta \theta^{\dagger}-I\right) W(s)$, the set $A=\{0\} \subset Y=\mathbb{S}_{+}^{d}$. This guarantees the existence of a measurable function $\tau_{0}: \mathcal{S} \rightarrow \Theta$ such that it minimizes pointwise $\varphi(s, \cdot)$ for any $s \in \mathcal{S}$ on the set $\left\{\theta \in \mathbb{S}_{+}^{\mathrm{d}} \mid \operatorname{Ran}(W(s)) \subset \operatorname{Ran}(\theta)\right\}$. The uniqueness of $\tau_{\rho}(s)$ for each $s \in \mathcal{S}$ guarantees that $\tau_{\rho}=\tau_{0}$ is measurable as desired.

## C Proof of Ex. 1

In this section we report the proof of the statement in Ex. 1.
Example 1 (Clusters). Let $\mathcal{S}=\mathbb{R}^{q}$ be the side information space, for some integer $\mathrm{q}>0$. Let $\rho$ be such that the side information marginal distribution $\rho_{\mathcal{S}}$ is given by a uniform mixture of $m$ uniform distributions. More precisely, let $\rho_{\mathcal{S}}=\frac{1}{m} \sum_{i=1}^{m} \rho_{\mathcal{S}}^{(i)}$, with $\rho_{\mathcal{S}}^{(i)}=\mathcal{U}\left(\mathcal{B}\left(\mathfrak{a}_{i}, 1 / 2\right)\right)$ the uniform distribution on the ball of radius $1 / 2$ centered at $a_{i} \in \mathcal{S}$, characterizing the cluster i. For a given side information s, a task $\mu \sim \rho(\cdot \mid s)$ is sampled such that: 1) its inputs' marginal $\eta_{\mu}$ is a distribution with constant covariance matrix $C(s)=\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} \mathbb{E}_{x \sim \eta_{\mu}} \chi x^{\top}=C$, for some $\left.\mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}}, 2\right) w_{\mu}$ is sampled from a distribution with conditional covariance matrix $W(s)=\mathbb{E}_{\mu \sim \rho(\cdot \mid s)} w_{\mu} w_{\mu}^{\top}$, with $W(s)$ such that $\left(\mathrm{C}^{1 / 2} W(s) \mathrm{C}^{1 / 2}\right)\left(\mathrm{C}^{1 / 2} \mathrm{~W}(\mathrm{p}) \mathrm{C}^{1 / 2}\right)=0$ if $\mathrm{s} \neq \mathrm{p}$. Then,

$$
\mathbb{E}_{s \sim \rho_{\mathcal{S}}}\left\|C(s)^{1 / 2} W(s)^{1 / 2}\right\|_{*}=\frac{1}{\sqrt{m}}\left\|W_{\rho}^{1 / 2} C_{\rho}^{1 / 2}\right\|_{*} .
$$

Proof. According to the setting described in the example, we can rewrite the following:

$$
\begin{align*}
\mathbb{E}_{s \sim \rho_{\mathcal{S}}}\left\|C(s)^{1 / 2} W(s)^{1 / 2}\right\|_{*} & =\mathbb{E}_{s \sim \rho_{\mathcal{S}}}\left\|C^{1 / 2} W(s)^{1 / 2}\right\|_{*} \\
& =\mathbb{E}_{s \sim \rho_{\mathcal{S}}} \operatorname{Tr}\left(\left(C^{1 / 2} W(s) C^{1 / 2}\right)^{1 / 2}\right) \\
& =\frac{1}{m} \sum_{i=1}^{m} \mathbb{E}_{s \sim \rho_{\mathcal{S}}} \operatorname{Tr}\left(\left(C^{1 / 2} W(s) C^{1 / 2}\right)^{1 / 2}\right) \\
& =\frac{1}{m} \sum_{i=1}^{m} \operatorname{Tr}\left(\left(C^{1 / 2} W\left(a_{i}\right) C^{1 / 2}\right)^{1 / 2}\right)  \tag{73}\\
& =\frac{1}{m} \operatorname{Tr}\left(\sum_{i=1}^{m}\left(C^{1 / 2} W\left(a_{i}\right) C^{1 / 2}\right)^{1 / 2}\right) \\
& =\frac{1}{m} \operatorname{Tr}\left(\left(\sum_{i=1}^{m} C^{1 / 2} W\left(a_{i}\right) C^{1 / 2}\right)^{1 / 2}\right),
\end{align*}
$$

where, in the first equality we have exploited the fact that $\mathrm{C}(s)$ is a constant matrix C , in the second equality we have applied the definition of the rewriting of the trace norm of a
matrix $A$ as $\|A\|_{*}=\operatorname{Tr}\left(\left(A A^{\top}\right)^{1 / 2}\right)$, in the third and fourth equality we have exploited the assumption on $\rho_{\mathcal{S}}$, and finally, in the last equality, by point 2 ), we managed to apply the fact that, for two matrices $A, B \in \mathbb{S}_{+}^{d}$ such that $A^{1 / 2} B^{1 / 2}=B^{1 / 2} A^{1 / 2}=0$, we have

$$
\begin{equation*}
\left(A^{1 / 2}+B^{1 / 2}\right)\left(A^{1 / 2}+B^{1 / 2}\right)=A+B \Longrightarrow(A+B)^{1 / 2}=A^{1 / 2}+B^{1 / 2} . \tag{74}
\end{equation*}
$$

On the other hand, we observe that we can also write the following:

$$
\begin{align*}
\left\|C_{\rho}^{1 / 2} W_{\rho}^{1 / 2}\right\|_{*} & =\left\|C^{1 / 2} W_{\rho}^{1 / 2}\right\|_{*} \\
& =\operatorname{Tr}\left(\left(C^{1 / 2} W_{\rho} C^{1 / 2}\right)^{1 / 2}\right) \\
& =\operatorname{Tr}\left(\left(C^{1 / 2} \mathbb{E}_{s \sim \rho_{\mathcal{S}}} W(s) C^{1 / 2}\right)^{1 / 2}\right) \\
& =\operatorname{Tr}\left(\left(C^{1 / 2} \frac{1}{m} \sum_{i=1}^{m} \mathbb{E}_{s \sim \rho_{\mathcal{S}}} W(s) C^{1 / 2}\right)^{1 / 2}\right)  \tag{75}\\
& =\frac{1}{\sqrt{m}} \operatorname{Tr}\left(\left(\sum_{i=1}^{m} C^{1 / 2} W\left(a_{i}\right) C^{1 / 2}\right)^{1 / 2}\right) \\
& =\frac{1}{\sqrt{m}} \operatorname{Tr}\left(\left(C^{1 / 2} \sum_{i=1}^{m} W\left(a_{i}\right) C^{1 / 2}\right)^{1 / 2}\right)
\end{align*}
$$

where, in the first equality we have exploited the fact that $C(s)$ is a constant matrix $C$, in the second equality we have applied the definition of the rewriting of the trace norm of a matrix $A$ as $\|A\|_{*}=\operatorname{Tr}\left(\left(A A^{\top}\right)^{1 / 2}\right)$ and in the fourth and fifth equality we have exploited the assumption on $\rho_{\mathcal{S}}$. The desired statement directly derives from combining Eq. (73) and Eq. (75).

## 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. 5. We start from proving the matricial rewriting of Prop. 3 in App. D.1. We then prove in App. D. 2 the properties of the surrogate functions in Prop. 4. Then, in App. D.3, we prove the convergence rate of Alg. 1 on the surrogate problem in Eq. (22).

## D. 1 Proof of Prop. 3

We start from proving the matricial rewriting of Prop. 3..
Proposition 3 (Matricial Re-formulation of $\tau_{M}(s)$ ). Let $\tau$ be as in Eq. (16). Then,

$$
\begin{equation*}
\tau(s)=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)^{\top}\right) \mathrm{H}_{\mathrm{M}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)+\mathrm{C} \tag{17}
\end{equation*}
$$

where $\mathrm{I}_{\mathrm{d}}$ is the identity in $\mathbb{R}^{\mathrm{d} \times \mathrm{d}}$ and $\mathrm{H}_{\mathrm{M}}$ is the matrix in $\mathbb{R}^{\mathrm{dk} \times \mathrm{dk}}$ defined by the entries

$$
\left(\mathrm{H}_{M}\right)_{(\mathrm{i}-1) \mathrm{k}+\mathrm{h},(\mathrm{j}-1) \mathrm{k}+z}=\langle\mathrm{M}(:, \mathfrak{i}, \mathrm{h}), \mathrm{M}(:, \mathfrak{j}, z)\rangle
$$

with $i, j=1, \ldots, d$ and $h, z=1, \ldots, k$.

Proof. We start from observing that for any $i, j=1, \ldots, d$, we can rewrite the following

$$
\begin{align*}
&\left((M \Phi(s))^{\top} M \Phi(s)\right)_{i, j}=\left\langle(M \Phi(s))^{\top}(\mathfrak{i},:),(M \Phi(s))(:, \mathfrak{j})\right\rangle \\
&=\langle(M \Phi(s))(:, \mathfrak{i}),(M \Phi(s))(:, \mathfrak{j})\rangle \\
&=\sum_{\mathbf{q}=1}^{m}(M \Phi(s))(:, \mathfrak{i})_{\mathfrak{q}}(M \Phi(s))(:, \mathfrak{j})_{\mathfrak{q}} \\
&=\sum_{\mathfrak{q}=1}^{m}\left(\sum_{h=1}^{k} M_{q, i, h} \Phi(s)_{h}\right)\left(\sum_{z=1}^{k} M_{q}, \mathfrak{j}, z\right. \\
& \Phi  \tag{76}\\
& \hline
\end{align*}
$$

We now observe that for any $i, j=1, \ldots, d$, we can rewrite the following

$$
\begin{align*}
& \left(\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right) \mathrm{H}_{M}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)\right)_{\mathrm{i}, \mathrm{j}}=\left\langle\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right)(\mathrm{i},:),\left(\mathrm{H}_{M}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)\right)(:, \mathfrak{j})\right\rangle \\
& =\left\langle\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)(:, \mathfrak{i}),\left(\mathrm{H}_{\mathrm{M}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)\right)(:, \mathfrak{j})\right\rangle \\
& =\sum_{n=1}^{k d}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)_{\mathrm{n}, \mathrm{i}}\left(\mathrm{H}_{\mathrm{M}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)\right)_{\mathrm{n}, \mathrm{j}} \\
& =\sum_{n=1}^{k d}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)_{\mathrm{n}, \mathrm{i}}\left\langle\mathrm{H}_{\mathrm{M}}(\mathrm{n},:),\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)(:, \mathfrak{j})\right\rangle \\
& =\sum_{n=1}^{k d}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right)_{n, i} \sum_{p=1}^{k d}\left(\mathrm{H}_{M}\right)_{n, p}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right)_{p, j} \\
& =\sum_{n=1}^{k d} \sum_{p=1}^{k d}\left(I_{d} \otimes \Phi(s)\right)_{n, i}\left(H_{M}\right)_{n, p}\left(I_{d} \otimes \Phi(s)\right)_{p, j} \\
& =\sum_{n=1}^{k d} \sum_{p=1}^{k d} \Phi(s)_{h} \delta_{n,(i-1) k+h}\left(H_{M}\right)_{n, p} \Phi(s)_{z} \delta_{\mathfrak{p},(j-1) k+z} \\
& =\sum_{h=1}^{k} \sum_{z=1}^{k} \Phi(s)_{h} \Phi(s)_{z}\left(\mathrm{H}_{M}\right)_{(i-1) k+h,(j-1) k+z}, \tag{77}
\end{align*}
$$

where, in the seventh equality we have exploited the fact that, by definition,

$$
\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right)_{n, i}=\left\{\begin{array}{ll}
\Phi(\mathrm{s})_{\mathrm{r}} & \text { if } \mathrm{r}=\mathrm{n}-(\mathrm{i}-1) \mathrm{k}  \tag{78}\\
0 & \text { otherwise }
\end{array}=\Phi(\mathrm{s})_{\mathrm{r}} \delta_{n, r+(i-1) k} .\right.
$$

and in the last equality we have defined the new indexes $h, z=1, \ldots, k$ as

$$
\begin{equation*}
h=n-(i-1) k \quad z=p-(j-1) k \tag{79}
\end{equation*}
$$

and, as consequence, we have rewritten

$$
\begin{equation*}
n=(i-1) k+h \quad p=(j-1) k+z . \tag{80}
\end{equation*}
$$

As, a consequence, if we define $H_{M}$ as the matrix in $\mathbb{R}^{d k \times d k}$ with entries

$$
\begin{equation*}
\left(\mathrm{H}_{\mathrm{M}}\right)_{(\mathrm{i}-1) \mathrm{k}+\mathrm{h},(\mathrm{j}-1) \mathrm{k}+z}=\langle\mathrm{M}(:, \mathfrak{i}, \mathrm{h}), \mathrm{M}(:, \mathfrak{j}, z)\rangle \tag{81}
\end{equation*}
$$

with $i, j=1, \ldots, d$ and $h, z=1, \ldots, k$, then, Eq. (76):

$$
\begin{equation*}
\left((M \Phi(s))^{\top} M \Phi(s)\right)_{i, j}=\sum_{h=1}^{k} \sum_{z=1}^{k} \Phi(s)_{h} \Phi(s)_{z}\langle M(:, i, h), M(:, j, z)\rangle \tag{82}
\end{equation*}
$$

and Eq. (77):

$$
\begin{equation*}
\left(\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right) \mathrm{H}_{\mathrm{M}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)\right)_{\mathrm{i}, \mathrm{j}}=\sum_{\mathrm{h}=1}^{\mathrm{k}} \sum_{z=1}^{\mathrm{k}} \Phi(\mathrm{~s})_{\mathrm{h}} \Phi(\mathrm{~s})_{z}\left(\mathrm{H}_{\mathrm{M}}\right)_{(\mathrm{i}-1) \mathrm{k}+\mathrm{h},(\mathfrak{j}-1) \mathrm{k}+z} \tag{83}
\end{equation*}
$$

coincide. This coincides with the first desired statement. In order to prove the statement $H_{M} \in \mathbb{S}_{+}^{d k}$, we show that $H_{M}=A_{M}^{\top} A_{M}$, where $A_{M}$ is the matrix in $\mathbb{R}^{m \times d k}$ defined as

$$
\begin{equation*}
A_{M}(:,(i-1) k+h)=M(:, i, h) . \tag{84}
\end{equation*}
$$

We start from recalling that, by definition of $\mathrm{H}_{M}$, we have

$$
\begin{equation*}
\left(\mathrm{H}_{\mathrm{M}}\right)_{(\mathrm{i}-1) \mathrm{k}+\mathrm{h},(j-1) \mathrm{k}+z}=\langle\mathrm{M}(:, \mathfrak{i}, \mathfrak{h}), M(:, \mathfrak{j}, z)\rangle . \tag{85}
\end{equation*}
$$

Moreover, we observe that, for any $p, q=1, \ldots, k d$,

$$
\begin{equation*}
\left(A_{M}^{\top} A_{M}\right)_{p, q}=\left\langle\left(A_{M}^{\top}\right)(p,:), A_{M}(:, q)\right\rangle_{\mathbb{R}^{m}}=\left\langle A_{M}(:, p), A_{M}(:, q)\right\rangle . \tag{86}
\end{equation*}
$$

As a consequence, the desired statement is satisfied if we define

$$
\begin{equation*}
\left(A_{M}\right)(:,(i-1) k+h)=M(:, i, h) . \tag{87}
\end{equation*}
$$

We now prove the last statement. Let $\left(e_{i}\right)_{i=1}^{d}$ be the canonical basis in $\mathbb{R}^{d}$. By the definition of the trace and the rewriting of $\tau(s)$ in Prop. 3, denoting by vec the vectorization operation,
we can rewrite

$$
\begin{align*}
\operatorname{Tr}(\tau(s)) & =\sum_{i=1}^{d}\left\langle e_{i}, \tau(s) e_{i}\right\rangle \\
& =\sum_{i=1}^{d}\left\langle e_{i},\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)^{\top}\right) \mathrm{H}_{M}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right) e_{i}\right\rangle \\
& =\sum_{i=1}^{\mathrm{d}} e_{i}^{\top}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right) \mathrm{H}_{M}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right) e_{i} \\
& =\sum_{i=1}^{\mathrm{d}}\left(\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right) e_{i}\right)^{\top} H_{M}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right) e_{i}  \tag{88}\\
& =\sum_{i=1}^{\mathrm{d}}\left(\operatorname{vec}\left(\Phi(\mathrm{~s}) e_{i}^{\top}\right)\right)^{\top} \mathrm{H}_{M} \operatorname{vec}\left(\Phi(\mathrm{~s}) e_{i}^{\top}\right) \\
& =\operatorname{Tr}\left(\mathrm{H}_{M} \sum_{i=1}^{d} \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right) \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right)^{\top}\right) \\
& \leq \operatorname{Tr}\left(\mathrm{H}_{M}\right)\left\|\sum_{i=1}^{d} \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right) \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right)^{\top}\right\|_{\infty} \\
& =\operatorname{Tr}\left(\mathrm{H}_{M}\right)\|\Phi(s)\|_{\mathbb{R}^{\mathrm{k}}}^{2},
\end{align*}
$$

where, in the fifth equality, we have applied the relation

$$
\begin{equation*}
\left(C^{\top} \otimes A\right) \operatorname{vec}(B)=\operatorname{vec}(A B C) \tag{89}
\end{equation*}
$$

with $A=\Phi(s), B=e_{i}^{\top}$ and $C=I_{d}$, i.e.

$$
\begin{equation*}
\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right) e_{\mathrm{i}}=\operatorname{vec}\left(\Phi(\mathrm{s}) e_{\mathrm{i}}^{\top}\right) \tag{90}
\end{equation*}
$$

in the inequality we have applied Holder's inequality and in the last equality we have applied the following proposition.

Proposition 11. For any $i=1, \ldots, d$, define

$$
\begin{equation*}
v_{i}=\operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right) \tag{91}
\end{equation*}
$$

Then,

$$
\begin{equation*}
\left\|\sum_{i=1}^{d} \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right) \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right)^{\top}\right\|_{\infty}=\left\|\sum_{i=1}^{d} v_{i} v_{i}^{\top}\right\|_{\infty}=\|\Phi(s)\|^{2} . \tag{92}
\end{equation*}
$$

Proof. We start from observing that, for any $i, j=1, \ldots, d$, we have

$$
\begin{align*}
v_{i}^{\top} v_{j} & =\operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right)^{\top} \operatorname{vec}\left(\Phi(s) e_{j}^{\top}\right) \\
& =\operatorname{Tr}\left(e_{i} \Phi(s)^{\top} \Phi(s) e_{j}^{\top}\right) \\
& =\operatorname{Tr}\left(\Phi(s)^{\top} \Phi(s) e_{j}^{\top} e_{i}\right)  \tag{93}\\
& =\Phi(s)^{\top} \Phi(s) e_{j}^{\top} e_{i} \\
& =\|\Phi(s)\|^{2} \delta_{i, j},
\end{align*}
$$

where, in the second equality, we have used the property of the operator vec:

$$
\begin{equation*}
\operatorname{vec}(A)^{\top} \operatorname{vec}(B)=\operatorname{Tr}\left(A^{\top} B\right) \tag{94}
\end{equation*}
$$

with

$$
\begin{equation*}
A=\Phi(s) e_{i}^{\top} \quad B=\Phi(s) e_{j}^{\top} . \tag{95}
\end{equation*}
$$

As a consequence, the vectors

$$
\begin{equation*}
\tilde{v}_{i}=\frac{v_{i}}{\left\|v_{i}\right\|}=\frac{v_{i}}{\|\Phi(s)\|} \quad i=1, \ldots, d \tag{96}
\end{equation*}
$$

form an orthonormal basis of the space. Moreover, we can rewrite the operator above as follows

$$
\begin{equation*}
\sum_{i=1}^{\mathrm{d}} \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right) \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right)^{\top}=\sum_{i=1}^{\mathrm{d}} v_{i} v_{i}^{\top}=\sum_{i=1}^{\mathrm{d}}\|\Phi(s)\|^{2} \tilde{v}_{i} \tilde{v}_{i}^{\top} . \tag{97}
\end{equation*}
$$

The rewriting above coincides with the eigenvalues' decomposition of the operator: the vectors $\tilde{v_{i}}$ are the eigenvectors with associated constant eigenvalues $\|\Phi(s)\|^{2}$. As a consequence, we can conclude that

$$
\begin{equation*}
\left\|\sum_{i=1}^{d} \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right) \operatorname{vec}\left(\Phi(s) e_{i}^{\top}\right)^{\top}\right\|_{\infty}=\|\Phi(s)\|^{2} \tag{98}
\end{equation*}
$$

## D. 2 Proof of Prop. 4

We now prove the properties of the surrogate functions in Prop. 4.
Proposition 4 (Properties of the Surrogate Meta-Loss $\mathcal{L}$ ). For any $\mathrm{Z} \in \mathcal{D}$ and $s \in \mathcal{S}$, the function $\mathcal{L}(\cdot, \cdot, s, Z)$ is convex and one of its subgradients is given, for any $\mathrm{H} \in \mathbb{S}_{+}^{d k}$ and $\mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}}$, by

$$
\begin{align*}
& \nabla \mathcal{L}(\mathrm{H}, \cdot, \mathrm{~s}, \mathrm{Z})(\mathrm{C})=\hat{\nabla}  \tag{23}\\
& \nabla \mathcal{L}(\cdot, \mathrm{C}, \mathrm{~s}, \mathrm{Z})(\mathrm{H})=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right) \hat{\nabla}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right)
\end{align*}
$$

where

$$
\hat{\nabla}=-\frac{\lambda}{2} \tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger} w_{\tau_{H, C}(s)} w_{\tau_{\mathrm{H}, \mathrm{C}}(s)}^{\top} \tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger}+\frac{2 \mathrm{~L}^{2} \mathrm{X}^{\top} X}{\mathrm{n}^{2}} .
$$

Moreover, under Asm. 1 and Asm. 3, we have

$$
\|\nabla \mathcal{L}(\cdot, \cdot, s, Z)(H, C)\|_{F} \leq\left(1+K^{2}\right)(L R)^{2}\left(\frac{1}{2}+\frac{2}{n}\right)
$$

Proof. We are interested in studying the properties of the surrogate function $\mathcal{L}(\cdot, \cdot, s, Z)$ : $\mathbb{S}_{+}^{\mathrm{dk}} \times \mathbb{S}_{+}^{\mathrm{d}} \rightarrow \mathbb{R}$ in Eq. (22). We start from observing that, such a function coincides with the
composition of the function

$$
\begin{align*}
& \theta \in \mathbb{S}_{+}^{\mathrm{d}} \mapsto \Delta(\theta, Z)=F(\theta, Z)+G(\theta, Z) \in \mathbb{R} \\
& F(\theta, Z)=\min _{w \in \mathbb{R}^{\mathrm{d}}} \mathcal{R}_{Z, \theta}(w) \quad \mathcal{R}_{Z, \theta}(w)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(\left\langle x_{i}, w\right\rangle, y_{i}\right)+\frac{\lambda}{2}\left\langle w, \theta^{\dagger} w\right\rangle+\iota_{\operatorname{Ran}(\theta)}(w) \\
& G(\theta, Z)=\frac{2 L^{2}}{n} \operatorname{Tr}\left(\theta \frac{X^{\top} X}{n}\right) . \tag{99}
\end{align*}
$$

with the linear transformation

$$
\begin{equation*}
s \in \mathcal{S} \mapsto \tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s})=\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right) \mathrm{H}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right)+\mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}} . \tag{100}
\end{equation*}
$$

In other words, for any $\mathrm{H} \in \mathbb{S}_{+}^{\mathrm{dk}}$ and $\mathrm{C} \in \mathbb{S}_{+}^{\mathrm{d}}$, we can write

$$
\begin{equation*}
\mathcal{L}(\mathrm{H}, \mathrm{C}, \mathrm{~s}, \mathrm{Z})=\Delta\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s}), \mathrm{Z}\right)=\mathrm{F}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s}), \mathrm{Z}\right)+\mathrm{G}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s}), \mathrm{Z}\right) \tag{101}
\end{equation*}
$$

We now observe that both the functions $F(\cdot, Z)$ and $G(\cdot, Z)$ are both convex $(F(\cdot, Z)$ is convex since it is the minimum of a jointly convex function see Denevi et al. (2019b) and $G(\cdot, Z)$ is a linear function). As a consequence, the function $\Delta(\cdot, Z)$ is convex over $\mathbb{S}_{+}^{d}$. This implies the convexity of the surrogate function $\mathcal{L}(\cdot, \cdot, s, Z)$ over $\mathbb{S}_{+}^{\mathrm{dk}} \times \mathbb{S}_{+}^{d}$ (composition of a convex function with a linear transformation). In order to get the closed form of the gradient in Eq. (23) we proceed in a similar way as in Denevi et al. (2020). More precisely, we start from recalling that, as already observed in Denevi et al. (2019b), thanks to strong duality in the within-task problem, for any $\theta \in \mathbb{S}_{+}^{d}$, we can rewrite

$$
\begin{equation*}
F(\theta, Z)=\min _{w \in \operatorname{Ran}(\theta)} \mathcal{R}_{Z, \theta}(w)=\max _{\alpha \in \mathbb{R}^{n}}\left\{-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 n^{2}} \operatorname{Tr}\left(\theta X^{\top} \alpha \alpha^{\top} X\right)\right\}, \tag{102}
\end{equation*}
$$

where, $\ell_{i}^{*}(\cdot)$ denotes the Fenchel conjugate of $\ell_{i}(\cdot)=\ell\left(\cdot, y_{i}\right)$ and $\alpha \in \mathbb{R}^{n}$ coicides with the dual variable. As a consequence, we can rewrite

$$
\begin{align*}
\Delta(\theta, Z) & =F(\theta, Z)+G(\theta, Z) \\
& =\max _{\alpha \in \mathbb{R}^{n}}\left\{-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)-\frac{1}{2 n^{2}} \operatorname{Tr}\left(\theta X^{\top} \alpha \alpha^{\top} X\right)\right\}+\frac{2 L^{2}}{n} \operatorname{Tr}\left(\theta \frac{X^{\top} X}{n}\right)  \tag{103}\\
& =\max _{\alpha \in \mathbb{R}^{n}}\left\{-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)+\operatorname{Tr}\left(\theta\left(-\frac{X^{\top} \alpha \alpha^{\top} X}{2 n^{2}}+\frac{2 L^{2} X^{\top} X}{n^{2}}\right)\right)\right\} .
\end{align*}
$$

As a consequence, we have

$$
\begin{align*}
& \Delta\left(\tau_{H, C}, Z\right)=\max _{\alpha \in \mathbb{R}^{n}}\left\{-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)+\operatorname{Tr}\left(\left(I_{d} \otimes \Phi(s)^{\top}\right) H\left(I_{d} \otimes \Phi(s)\right)\left(-\frac{X^{\top} \alpha \alpha^{\top} X}{2 n^{2}}+\frac{2 L^{2} X^{\top} X}{n^{2}}\right)\right)\right. \\
& \left.+\operatorname{Tr}\left(C\left(-\frac{X^{\top} \alpha \alpha^{\top} X}{2 n^{2}}+\frac{2 L^{2} X^{\top} X}{n^{2}}\right)\right)\right\} \\
& =\max _{\alpha \in \mathbb{R}^{n}}\left\{-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)+\operatorname{Tr}\left(H\left(I_{d} \otimes \Phi(s)\right)\left(-\frac{X^{\top} \alpha \alpha^{\top} X}{2 n^{2}}+\frac{2 L^{2} X^{\top} X}{n^{2}}\right)\left(I_{d} \otimes \Phi(s)^{\top}\right)\right)\right. \\
& \left.+\operatorname{Tr}\left(C\left(-\frac{X^{\top} \alpha \alpha^{\top} X}{2 n^{2}}+\frac{2 L^{2} X^{\top} X}{n^{2}}\right)\right)\right\} \\
& =\max _{\alpha \in \mathbb{R}^{n}} \mathrm{Q}(\alpha, \mathrm{H}, \mathrm{C}, s, \mathrm{Z}) \text {, } \tag{104}
\end{align*}
$$

where we have introduced the function

$$
\begin{align*}
\mathrm{Q}(\alpha, \mathrm{H}, \mathrm{C}, \mathrm{~s}, \mathrm{Z})= & -\frac{1}{n} \sum_{i=1}^{n} \ell_{i}^{*}\left(\alpha_{i}\right)+\operatorname{Tr}\left(\mathrm{H}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right)\left(-\frac{X^{\top} \alpha \alpha^{\top} X}{2 \mathrm{n}^{2}}+\frac{2 \mathrm{~L}^{2} X^{\top} X}{\mathrm{n}^{2}}\right)\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)^{\top}\right)\right) \\
& +\operatorname{Tr}\left(\mathrm{C}\left(-\frac{X^{\top} \alpha \alpha^{\top} X}{2 \mathrm{n}^{2}}+\frac{2 \mathrm{~L}^{2} X^{\top} X}{n^{2}}\right)\right) . \tag{105}
\end{align*}
$$

Hence, applying (Denevi et al., 2019b, Lemma 44), we know that, once computed a maximizer $\alpha_{\tau_{\mu, C}(s)}$ of the function above $\alpha \in \mathbb{R}^{n} \mapsto \mathrm{Q}(\alpha, H, C, s, Z)$,

$$
\begin{equation*}
\nabla \mathrm{Q}\left(\alpha_{\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s})}, \cdot, \cdot, \mathrm{s}, \mathrm{Z}\right)(\mathrm{H}, \mathrm{C}) \in \frac{\partial \Delta\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s}), \mathrm{Z}\right)}{\partial(\mathrm{H}, \mathrm{C})}=\frac{\partial \mathcal{L}(\mathrm{H}, \mathrm{C}, \mathrm{~s}, \mathrm{Z})}{\partial(\mathrm{H}, \mathrm{C})} . \tag{106}
\end{equation*}
$$

As a consequence, since for a given matrix $A, \nabla \operatorname{Tr}(\cdot A)(H)=A$, we get that

$$
\begin{equation*}
\nabla \mathcal{L}(\cdot, \cdot, s, \mathrm{Z})(\mathrm{H}, \mathrm{C})=\left(\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})\right) \stackrel{\nabla}{\mathrm{D}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right), \stackrel{\nabla}{\nabla}\right) \in \frac{\partial \mathcal{L}(\mathrm{H}, \mathrm{C}, \mathrm{~s}, \mathrm{Z})}{\partial(\mathrm{H}, \mathrm{C})} \tag{107}
\end{equation*}
$$

with

$$
\begin{equation*}
\hat{\nabla}=-\frac{X^{\top} \alpha_{\tau_{\mathrm{H}, \mathrm{C}}(s)} \alpha_{\tau_{H, C}(s)}^{\top} X}{2 n^{2}}+\frac{2 \mathrm{~L}^{2} X^{\top} X}{n^{2}} . \tag{108}
\end{equation*}
$$

Finally, in order to get the desired closed form in Eq. (23), we just need to observe that, according to the optimality conditions of the within-task problem in (see (Denevi et al., 2019b, Lemma 44)) with $\theta \in \mathbb{S}_{+}^{d}$, we have that

$$
\begin{equation*}
\mathrm{X}^{\top} \alpha_{\theta}=-\mathrm{n} \theta^{\dagger} w_{\theta} . \tag{109}
\end{equation*}
$$

As a consequence, we can rewrite Eq. (108) as follows by using the primal solution of the within-task problem:

$$
\begin{equation*}
\hat{\nabla}=-\frac{\lambda}{2} \tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger} w_{\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s})} w_{\tau_{\mathrm{H}, \mathrm{C}}(s)}^{\top} \tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger}+\frac{2 \mathrm{~L}^{2} X^{\top} X}{\mathrm{n}^{2}} . \tag{110}
\end{equation*}
$$

Finally, we observe that, by the closed form in Eq. (23),

$$
\begin{equation*}
\|\nabla \mathcal{L}(\cdot, \cdot, s, Z)(\mathrm{H}, \mathrm{C})\|_{\mathrm{F}} \leq\|\nabla \mathcal{L}(\cdot, \cdot, s, \mathrm{Z})(\mathrm{H}, \mathrm{C})\|_{*} \leq \mathrm{A}+\mathrm{B}+\mathrm{C}+\mathrm{D} \tag{111}
\end{equation*}
$$

with

$$
\begin{align*}
& A=\left\|\left(I_{d} \otimes \Phi(s)\right) \frac{X^{\top} \alpha_{\tau_{H, C}(s)} \alpha_{\tau_{H, C}(s)}^{\top} X}{2 n^{2}}\left(I_{d} \otimes \Phi(s)^{\top}\right)\right\|_{*} \\
& B=\left\|\left(I_{d} \otimes \Phi(s)\right) \frac{2 L^{2} X^{\top} X}{n^{2}}\left(I_{d} \otimes \Phi(s)^{\top}\right)\right\|_{*} \\
& C=\left\|\frac{X^{\top} \alpha_{\tau_{H}, \mathrm{C}(s)} \alpha_{\tau_{H, C}(s)}^{\top} X}{2 n^{2}}\right\|_{*}  \tag{112}\\
& D=\left\|\frac{2 L^{2} X^{\top} X}{n^{2}}\right\|_{*} .
\end{align*}
$$

We now observe that all the matrices inside the trace norms above are positive semidefinite (as a matter of fact, if a matrix Q is positive semidefinite, then, $\mathrm{P}^{\top} \mathrm{QP}$ is positive semidefinite for any matrix P). As a consequence, all the trace norms above coincide with the trace of the corresponding matrices, namely,

$$
\begin{align*}
& \mathrm{A}=\operatorname{Tr}\left(\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right) \frac{\mathrm{X}^{\top} \alpha_{\tau_{\mathrm{H}, \mathrm{C}(s)}} \alpha_{\tau_{\mathrm{H}, \mathrm{C}(\mathrm{~s})}^{\top}} \mathrm{X}}{2 \mathrm{n}^{2}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(\mathrm{~s})^{\top}\right)\right) \\
& \mathrm{B}=\operatorname{Tr}\left(\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)\right) \frac{2 \mathrm{~L}^{2} \mathrm{X}^{\top} \mathrm{X}}{\mathrm{n}^{2}}\left(\mathrm{I}_{\mathrm{d}} \otimes \Phi(s)^{\top}\right)\right) \\
& \mathrm{C}=\operatorname{Tr}\left(\frac{\mathrm{X}^{\top} \alpha_{\tau_{\mathrm{H}, \mathrm{C}(s)}} \alpha_{\tau_{\mathrm{H}, \mathrm{C}(s)}} \mathrm{X}}{2 \mathrm{n}^{2}}\right)  \tag{113}\\
& \mathrm{D}=\operatorname{Tr}\left(\frac{2 \mathrm{~L}^{2} X^{\top} X}{\mathrm{n}^{2}}\right) .
\end{align*}
$$

We now observe that, proceeding as above in Eq. (88) and exploiting Asm. 3, we can write

$$
\begin{align*}
& \mathrm{A} \leq\|\Phi(\mathrm{s})\|^{2} \operatorname{Tr}\left(\frac{\left.\mathrm{X}^{\top} \alpha_{\tau_{\mathrm{H}, \mathrm{C}(\mathrm{c})} \alpha_{\tau_{H, C}(s)}^{\top} \mathrm{X}}^{2 \mathrm{n}^{2}}\right)=\|\Phi(s)\|^{2} \mathrm{C} \leq \mathrm{K}^{2} \mathrm{C}}{\mathrm{~B} \leq\|\Phi(\mathrm{s})\|^{2} \operatorname{Tr}\left(\frac{2 \mathrm{~L}^{2} X^{\top} \mathrm{X}}{n^{2}}\right)=\|\Phi(s)\|^{2} \mathrm{D} \leq \mathrm{K}^{2} \mathrm{D} .}\right.
\end{align*}
$$

Hence, combining everything in Eq. (111), we get

$$
\begin{equation*}
\|\nabla \mathcal{L}(\cdot, \cdot, s, Z)(H, C)\|_{F} \leq\left(1+\mathrm{K}^{2}\right)(\mathrm{C}+\mathrm{D}) \tag{115}
\end{equation*}
$$

The desired statement derives from observing that, since, by Asm. $1, \operatorname{Tr}\left(X^{\top} \alpha_{\tau_{\mu, C}(s)} \alpha_{\tau_{H}, \mathrm{C}(s)}^{\top} X\right) \leq$ $(n L R)^{2}$ (see (Denevi et al., 2019b, Lemma 44)) and $\operatorname{Tr}\left(X^{\top} X\right)=\operatorname{Tr}\left(X X^{\top}\right)=\sum_{i=1}^{n}\left\|x_{i}\right\|^{2} \leq$ $n R^{2}$, then

$$
\begin{equation*}
\mathrm{C} \leq \frac{(\mathrm{LR})^{2}}{2 \lambda} \quad \mathrm{D} \leq \frac{2(\mathrm{LR})^{2}}{\mathrm{n}} . \tag{116}
\end{equation*}
$$

## D. 3 Convergence rate of Alg. 1 on the surrogate problem in Eq. (22)

We now give the convergence rate of Alg. 1 on the surrogate problem in Eq. (22).
Proposition 12 (Convergence rate on the surrogate problem in Eq. (22)). Let $\overline{\mathrm{H}}$ and $\overline{\mathrm{C}}$ be the average of the iterations obtained from the application of Alg. 1 over the training data $\left(Z_{t}, s_{t}\right)_{t=1}^{\top}$ with constant meta-step size $\gamma>0$. Then, under Asm. 1 and Asm. 3, for any $\tau_{\mathrm{H}, \mathrm{C}} \in \mathcal{T}_{\Phi}$, in expectation w.r.t. the sampling of $\left(\mathrm{Z}_{\mathrm{t}}, s_{t}\right)_{\mathrm{t}=1}^{\top}$,

$$
\begin{equation*}
\mathbb{E} \hat{\mathcal{E}}_{\rho}\left(\tau_{\bar{H}, \bar{C}}\right)-\hat{\mathcal{E}}_{\rho}\left(\tau_{H, C}\right) \leq \frac{\gamma\left(1+\mathrm{K}^{2}\right)^{2}(\mathrm{LR})^{4}}{2 \lambda^{2}}\left(\frac{1}{2}+\frac{2}{n}\right)^{2}+\frac{\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}^{2}}{2 \gamma \mathrm{~T}} \tag{117}
\end{equation*}
$$

Proof. We observe that Alg. 1 coincides with projected Stochastic Gradient Descent applied to the convex and Lipschitz (see Prop. 4) surrogate problem in Eq. (22):

$$
\begin{equation*}
\min _{H \in \mathbb{S}_{+}^{\alpha k}, \mathrm{C} \in \mathbb{S}_{+}^{d}} \hat{\mathcal{E}}_{\rho}\left(\tau_{H, \mathrm{C}}\right) \quad \hat{\mathcal{E}}_{\rho}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)=\mathbb{E}_{(\mu, \mathrm{s}) \sim \rho} \mathbb{E}_{\mathrm{Z} \sim \mu^{n}} \mathcal{L}(\mathrm{H}, \mathrm{C}, \mathrm{~s}, \mathrm{Z}) \tag{118}
\end{equation*}
$$

As a consequence, by standard arguments (see e.g. (Shalev-Shwartz and Ben-David, 2014, Lemma 14.1, Thm. 14.8) and references therein), for any $\tau_{H, C} \in \mathcal{T}_{\Phi}$, we have

$$
\begin{equation*}
\mathbb{E} \hat{\mathcal{E}}_{\rho}\left(\tau_{\bar{H}, \bar{C}}\right)-\hat{\mathcal{E}}_{\rho}\left(\tau_{H, C}\right) \leq \frac{\gamma}{2 T} \sum_{\mathrm{t}=1}^{\mathrm{T}} \mathbb{E}\left\|\nabla \mathcal{L}\left(\cdot, \cdot, s, \mathrm{Z}_{\mathrm{t}}\right)\left(\mathrm{H}_{\mathrm{t}}, \mathrm{C}_{\mathrm{t}}\right)\right\|_{\mathrm{F}}^{2}+\frac{\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}^{2}}{2 \gamma \mathrm{~T}} . \tag{119}
\end{equation*}
$$

The desired statement derives from combining this bound with the bound on the norm of the meta-subgradients in Prop. 4.

## D. 4 Proof of Thm. 5

We now have all the ingredients necessary to prove Thm. 5 .
Theorem 5 (Excess Risk Bound for the Conditioning Function Returned by Alg. 1). Let Asm. 1 and Asm. 3 hold. For any $s \sim \rho_{\mathcal{S}}$, recall the conditional covariance matrices $W(s)$ and $\mathrm{C}(\mathrm{s})$ introduced in Thm. 1. Let $\tau_{\mathrm{H}, \mathrm{C}}$ be a fixed function in $\mathcal{T}_{\Phi}$ such that $\operatorname{Ran}(\mathrm{W}(\mathrm{s})) \subseteq$ $\operatorname{Ran}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{s})\right)$ for any $\mathrm{s} \sim \rho_{\mathcal{S}}$. Let $\overline{\mathrm{H}}$ and $\overline{\mathrm{C}}$ be the outputs of Alg. 1 applied to a sequence $\left(Z_{t}, s_{t}\right)_{t=1}^{\top}$ of i.i.d. pairs sampled from $\rho$ with meta-step size

$$
\begin{equation*}
\gamma=\frac{\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}}{\left(1+\mathrm{K}^{2}\right)(\mathrm{LR})^{2}}\left(\frac{1}{2}+\frac{2}{n}\right)^{-1} \frac{1}{\sqrt{T}} \tag{24}
\end{equation*}
$$

Then, in expectation w.r.t. the sampling of $\left(Z_{t}, s_{t}\right)_{t=1}^{\top}$,

$$
\begin{aligned}
\mathbb{E} \mathcal{E}_{\rho}\left(\tau_{\bar{H}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \leq & \frac{\mathbb{E}_{s \sim \rho_{\mathcal{S}}} \operatorname{Tr}\left(\tau_{\mathrm{H}, \mathrm{C}}(s)^{\dagger} \mathrm{W}(\mathrm{~s})\right)}{2}+\frac{2 \mathrm{~L}^{2} \mathbb{E}_{s \sim \rho_{\mathcal{S}}} \operatorname{Tr}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s}) \mathrm{C}(s)\right)}{n} \\
& +\left(\frac{1}{2}+\frac{2}{n}\right) \frac{\left(1+\mathrm{K}^{2}\right)(\mathrm{LR})^{2}\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}}{\sqrt{T}} .
\end{aligned}
$$

Proof. We start from observing thta, in expectation w.r.t. the meta-training set, for any fixed conditioning function $\tau_{H, C} \in \mathcal{T}_{\Phi}$, we can write the following decomposition

$$
\begin{align*}
\mathbb{E} \mathcal{E}_{\rho}\left(\tau_{\overline{\mathrm{H}}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} & \leq \mathbb{E} \hat{\mathcal{E}}_{\rho}\left(\tau_{\bar{H}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \\
& =\mathbb{E} \hat{\mathcal{E}}_{\rho}\left(\tau_{\overline{\mathrm{H}}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \pm \hat{\mathcal{E}}_{\rho}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)  \tag{120}\\
& =\underbrace{\mathbb{E} \hat{\mathcal{E}}_{\rho}\left(\tau_{\bar{H}, \overline{\mathrm{C}}}\right)-\hat{\mathcal{E}}_{\rho}\left(\tau_{H, \mathrm{C}}\right)}_{\mathrm{A}\left(\tau_{\mu, \mathrm{C}}\right)}+\underbrace{\hat{\mathcal{E}}_{\rho}\left(\tau_{H, \mathrm{C}}\right)-\mathcal{E}_{\rho}^{*}}_{\mathrm{B}\left(\tau_{H, \mathrm{C}}\right)},
\end{align*}
$$

where in the inequality above we have exploited the fact that, for any $\tau \in \mathcal{T}, \mathcal{E}_{\rho}(\tau) \leq \mathcal{\mathcal { E }}_{\rho}(\tau)$ (see Eq. (21)). We now observe that the term $\mathrm{A}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)$ can be controlled according to the convergence properties of the meta-algorithm in Alg. 1 as described in Prop. 12:

$$
\begin{equation*}
\mathbb{E} \hat{\mathcal{E}}_{\rho}\left(\tau_{\bar{H}, \overline{\mathrm{C}}}\right)-\hat{\mathcal{E}}_{\rho}\left(\tau_{H, C}\right) \leq \frac{\gamma\left(1+\mathrm{K}^{2}\right)^{2}(\mathrm{LR})^{4}}{2}\left(\frac{1}{2}+\frac{2}{n}\right)^{2}+\frac{\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}^{2}}{2 \gamma \mathrm{~T}} \tag{121}
\end{equation*}
$$

Regarding the term $\mathrm{B}\left(\tau_{\mathrm{H}, \mathrm{C}}\right)$, we observe that, for any $\tau$, we can rewrite

$$
\begin{align*}
\mathrm{B}(\tau) & =\hat{\mathcal{E}}_{\rho}(\tau)-\mathcal{E}_{\rho}^{*} \\
& =\mathbb{E}_{(\mu, s) \sim \rho} \mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}(A(\tau(s), Z))-\mathcal{R}_{\mu}\left(w_{\mu}\right)\right]+\frac{2 \mathrm{~L}^{2} \mathbb{E}_{(\mu, s) \sim \rho} \operatorname{Tr}\left(\tau(s) \mathbb{E}_{\chi \sim \eta_{\mu}} x x^{\top}\right)}{n} \\
& \leq \frac{\mathbb{E}_{(\mu, s) \sim \rho} \operatorname{Tr}\left(\tau(s)^{\dagger} w_{\mu} w_{\mu}^{\top}\right)}{2}+\frac{2 \mathrm{~L}^{2} \mathbb{E}_{(\mu, s) \sim \rho} \operatorname{Tr}\left(\tau(s) \mathbb{E}_{x \sim \eta_{\mu}} \chi x^{\top}\right)}{n}, \tag{122}
\end{align*}
$$

where in the inequality we have exploited the fact that, thanks to the definition of the algorithm, for any $(\mu, s) \sim \rho$, we can write

$$
\begin{equation*}
\mathbb{E}_{Z \sim \mu^{n}}\left[\mathcal{R}_{Z, \tau(s)}(\mathrm{A}(\tau(s), Z))-\mathcal{R}_{\mu}\left(w_{\mu}\right)\right] \leq \frac{\operatorname{Tr}\left(\tau(s)^{\dagger} w_{\mu} w_{\mu}^{\top}\right)}{2} . \tag{123}
\end{equation*}
$$

Combining the bounds on the two terms above in Eq. (120), we get

$$
\begin{align*}
\mathbb{E} \mathcal{E}_{\rho}\left(\tau_{\bar{H}, \overline{\mathrm{C}}}\right)-\mathcal{E}_{\rho}^{*} \leq & \frac{\mathbb{E}_{(\mu, \mathrm{s}) \sim \rho} \operatorname{Tr}\left(\tau_{H, \mathrm{C}}(\mathrm{~s})^{\dagger} w_{\mu} w_{\mu}^{\top}\right)}{2}+\frac{2 \mathrm{~L}^{2} \mathbb{E}_{(\mu, \mathrm{s}) \sim \rho} \operatorname{Tr}\left(\tau_{\mathrm{H}, \mathrm{C}}(\mathrm{~s}) \mathbb{E}_{x \sim \eta_{\mu}} x x^{\top}\right)}{n} \\
& +\frac{\gamma\left(1+\mathrm{K}^{2}\right)^{2}(\mathrm{LR})^{4}}{2}\left(\frac{1}{2}+\frac{2}{n}\right)^{2}+\frac{\left\|\left(\mathrm{H}-\mathrm{H}_{0}, \mathrm{C}-\mathrm{C}_{0}\right)\right\|_{\mathrm{F}}^{2}}{2 \gamma \top} . \tag{124}
\end{align*}
$$

The desired statement derives from optimizing w.r.t. the hyper-parameter $\gamma>0$.

## E Experimental Details

In this section we report the experimental details we missed in the main body. Specifically, we report the details regarding the tuning of the hyper-parameter $\gamma$ and the characteristics of the machine we used for running our experiments.

Synthetic Clusters. In order to tune the hyper-parameter $\gamma$ we applied the procedure
above with 14 candidates values for $\gamma$ in the range $\left[10^{-5}, 10^{5}\right]$ with logarithmic spacing and we evaluated the performance of the estimated meta-parameters (linear representations) by using $T=T_{t r}=500, T_{\text {va }}=300, T_{\text {te }}=100$ 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.

Lenk Dataset. In order to tune the hyper-parameter $\gamma$ we applied the procedure above with 14 candidates values for $\gamma$ in the range $\left[10^{-5}, 10^{5}\right]$ with logarithmic spacing and we evaluated the performance of the estimated meta-parameters (linear representations) by using $T=T_{\text {tr }}=100, \mathrm{~T}_{\text {va }}=40, \mathrm{~T}_{\text {te }}=30$ of the available tasks for meta-training, metavalidation and meta-testing, respectively. In order to train and to test the inner algorithm, we splitted each within-task dataset into $n=n_{t r}=16$ for training and $n_{t e}=4$ for test.

Movieles-100k Dataset. In order to tune the hyper-parameter $\gamma$ we applied the procedure above with 14 candidates values for $\gamma$ in the range $\left[10^{-5}, 10^{5}\right]$ with logarithmic spacing and we evaluated the performance of the estimated meta-parameters (linear representations) by using $T=T_{\text {tr }}=200, T_{\text {va }}=100, T_{\text {te }}=100$ 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 $n=n_{\text {tr }}=15$ for training and $n_{\text {te }}=5$ for test.

Jester-1 Dataset. In order to tune the hyper-parameter $\gamma$ we applied the procedure above with 14 candidates values for $\gamma$ in the range $\left[10^{-5}, 10^{5}\right]$ with logarithmic spacing and we evaluated the performance of the estimated meta-parameters (linear representations) by using $\mathrm{T}=\mathrm{T}_{\mathrm{tr}}=250, \mathrm{~T}_{\mathrm{va}}=100, \mathrm{~T}_{\mathrm{te}}=100$ of the available tasks for meta-training, metavalidation and meta-testing, respectively. In order to train and to test the inner algorithm, we splitted each within-task dataset into $\mathfrak{n}=\mathfrak{n}_{\text {tr }}=15$ for training and $n_{t e}=5$ for test.

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


[^0]:    ${ }^{1}$ Department of Computer Science, University College London, London, United Kingdom
    ${ }^{2}$ Computational Statistics and Machine Learning, Istituto Italiano di Tecnologia, Genova, Italy

