
Multitask Learning with No Regret: from Improved Confidence Bounds to Active Learning

Pier Giuseppe Sessa*
ETH Zürich
piergiuseppe.sessa@inf.ethz.ch

Pierre Laforgue*
Università degli Studi di Milano
pierre.laforgue@unimi.it

Nicolò Cesa-Bianchi
Università degli Studi di Milano
Politecnico di Milano
nicolo.cesa-bianchi@unimi.it

Andreas Krause
ETH Zürich
krausea@ethz.ch

Abstract

Multitask learning is a powerful framework that enables one to simultaneously learn multiple related tasks by sharing information between them. Quantifying uncertainty in the estimated tasks is of pivotal importance for many downstream applications, such as online or active learning. In this work, we provide novel confidence intervals for multitask regression in the challenging agnostic setting, i.e., when neither the similarity between tasks nor the tasks' features are available to the learner. The obtained intervals do not require i.i.d. data and can be directly applied to bound the regret in online learning. Through a refined analysis of the multitask information gain, we obtain new regret guarantees that, depending on a task similarity parameter, can significantly improve over treating tasks independently. We further propose a novel online learning algorithm that achieves such improved regret without knowing this parameter in advance, i.e., automatically adapting to task similarity. As a second key application of our results, we introduce a novel multitask active learning setup where several tasks must be simultaneously optimized, but only one of them can be queried for feedback by the learner at each round. For this problem, we design a no-regret algorithm that uses our confidence intervals to decide which task should be queried. Finally, we empirically validate our bounds and algorithms on synthetic and real-world (drug discovery) data.

1 Introduction

In many real-world applications, one often faces multiple related tasks to be solved sequentially or simultaneously. The goal of multitask learning (MTL) [4] is to leverage the similarities across the tasks to obtain more accurate and robust models. Indeed, by jointly learning multiple tasks, MTL can exploit their statistical dependencies, yielding better generalization and faster learning than treating each task independently. MTL has gained significant attention in recent years, as it has been shown to be effective in a wide range of applications, including natural language processing, computer vision, federated learning, and drug discovery, see e.g., [11, 19, 16, 27].

A very natural model for learning across multiple tasks is the agnostic multitask (MT) regression approach of [13]. This utilizes a multitask kernel that can interpolate between running N (number of tasks) independent regressions, and regressing all tasks to their common average, depending on a tunable parameter. Notably, such a kernel does not require any knowledge neither about tasks' features

*equal contribution

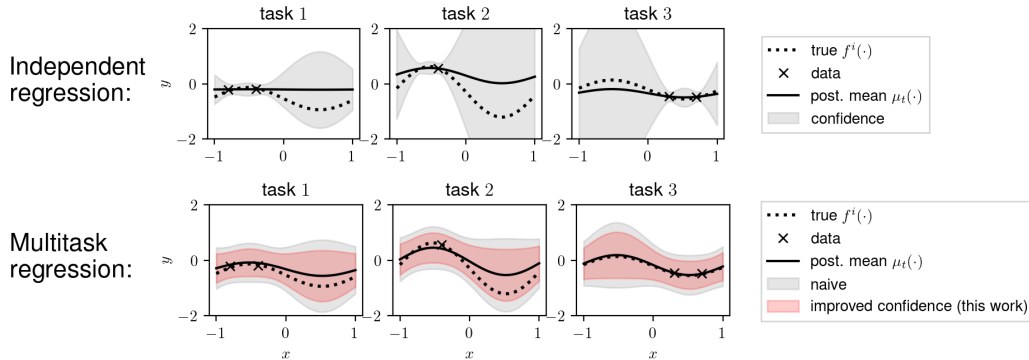


Figure 1: Independent vs. Multitask (MT) regression. MT regression leverages data coming from multiple related tasks and can yield more accurate and more confident estimates. In this work, we show naive confidence intervals are overly conservative and provide improved ones (shaded in red).

nor about their similarity, thus finding good application in several domains. For instance, Cavallanti et al. [5] study it for online classification, and Cesa-Bianchi et al. [6] for online convex optimization.

However, it is much less understood how to *quantify the uncertainty* of such MT regression, i.e., assessing confidence in the estimated tasks. In particular, as also outlined by [13] as an open problem, it is important to assess their generalization error as a function of the kernel parameter. Appropriately characterizing these confidence intervals is indeed of crucial importance for a whole set of downstream applications. More concretely, multitask confidence intervals are used in online learning to inform the next decision to be made [6]. In active learning—as we show next—these intervals are pivotal to deciding the most informative task to query.

In this work, we study the agnostic MT regression setup of [13], and provide *new multitask confidence intervals* (see Figure 1 for a visualization) for the full range of the kernel parameter. Our intervals hold in the so-called adaptive setting, i.e., without requiring i.i.d. data, and are *tighter up to a \sqrt{N} factor* than the naive ones employed in [6]. Moreover, we provide the first bounds for the information gain of MT regression and utilize them—together with the derived intervals—to obtain *tighter online learning guarantees*. The latter depend on a task similarity parameter and can significantly improve over treating tasks independently. Additionally, we propose an adaptive no-regret algorithm that exploits task similarity without knowing this parameter in advance. Finally, we consider a novel multitask *active learning* setup, where tasks should be simultaneously optimized but only one of them can be queried at each round. We show that the newly derived intervals are also crucial in such a setting, and provide a new algorithm that ensures sublinear regret. We demonstrate the superiority of the derived intervals over previously proposed algorithms on synthetic as well as real-world drug discovery tasks.

Related work. The agnostic MT regression approach of [13] reduces the learning of N tasks to a single regression problem, as a function of the MT kernel parameter. When combined with support vector machines, it was shown effective in a series of classification problems [13, 24], and since then was studied in various further settings. Cavallanti et al. [5], e.g., analyze mistake bounds for online MT classification algorithms as a function of the kernel parameter. Cesa-Bianchi et al. [6], instead, utilize the MT kernel to prove regret bounds in online MT learning with bandit feedback. Inspired by this, [10] focuses on learning more general kernel structures from data. An important question not addressed by previous work, though, is how to properly quantify the uncertainty of the obtained task estimates. This problem is well-understood in single-task learning (e.g., [2, 25, 8]) but remains largely unexplored in MT domains. As shown in [6], MT confidence intervals can in principle be obtained by a naive application of the single-task guarantees of [2]. However, as we show in Section 2, the so-obtained intervals are extremely conservative and—as a result—can hamper the MT learning performance. Our intervals are tighter by a factor up to \sqrt{N} w.r.t. the naive ones from [6], yielding novel online learning regret guarantees which can provably improve over treating tasks independently.

Compared to MT online learning [5, 6], where a single task is revealed to the learner at each round, a series of works have considered learning multiple tasks *simultaneously*, i.e., taking a decision for each one of them. Dekel et al. [12], e.g., propose the use of a shared loss function to account for tasks’ relatedness, Lugosi et al. [20] studies the computational tractability of taking multiple actions with joint constraints, while Cavallanti et al. [5] propose a matrix-based extension of the multitask Perceptron algorithm. In all of these works, however, the learner receives feedback from *all* the

tasks. In Section 4, instead, we focus on the challenging setup where only one task can be queried by the learner at each round. Hence, in addition to choosing good actions, the learner faces the *active learning* task of assessing the most informative feedback, in order to achieve sublinear regret. Perhaps most related to ours, is the offline contextual Bayesian optimization setup of [7, 18] where the goal is to compute the best strategy for each context (task) with minimal function interactions. However, unlike us, [7, 18] do not guarantee sublinear regret but provide only sample-complexity results.

Finally, we note that MT confidence intervals and regret guarantees were also recently derived by [9], albeit in a different setup and regression model. Indeed, the authors of [9] focus on multi-objective optimization where they ought to learn *multi-output* functions (each output corresponding to a task) using matrix-valued kernels. Although their setup can be related to ours, it crucially requires all tasks to be observed at each round, leading to different challenges than ours, see Appendix A.1 for details.

Notation. We use $[N] := \{1, \dots, N\}$, $\mathbb{1}_N$ for the vector in \mathbb{R}^N full of ones. Norms of functions are always taken w.r.t. the natural RKHS norm, so that we drop the subscript for simplicity of writing.

2 Improved Confidence Intervals for Multitask Kernel Regression

In this section, we introduce the MT kernel regression setting, and prove our refined confidence intervals. Of independent interest, these results are then leveraged in Sections 3 and 4 to derive novel regret bounds for online and active multitask learning. All proofs are deferred to the Appendix.

2.1 Multitask Kernel Regression

Given an input space \mathcal{X} , equipped with a (single task) scalar kernel $k_{\mathcal{X}}: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, the goal of MT kernel regression is to jointly learn N different functions f_1, \dots, f_N from \mathcal{X} to \mathbb{R} , all belonging to $\mathcal{H}_{k_{\mathcal{X}}}$, the RKHS associated to $k_{\mathcal{X}}$. To do so, the learner is given a set of triplets $\{(i_s, x_s), y_s\}_{s=1}^t$ consisting of a measured task index $i_s \in [N]$, a measured point $x_s \in \mathcal{X}$, and a noisy measurement $y_s = f_{i_s}(x_s) + \xi_s$, where ξ_s is an independent random variable to be specified later. We can further define the multitask function $f^{\text{mt}}: [N] \times \mathcal{X} \rightarrow \mathbb{R}$ such that $f^{\text{mt}}(i, \cdot) = f_i$, and the multitask kernel

$$k((i, x), (i', x')) = k_{\mathcal{T}}(i, i') \cdot k_{\mathcal{X}}(x, x'), \quad (1)$$

where $k_{\mathcal{T}}$ is a kernel on the tasks. In certain cases, the latter might be given as input to the learner, either under the form of the task Gram matrix, or via task features and an assumed (e.g., linear) similarity [17]. However, in practice such information is usually not accessible to the learner. In such a case, a standard *agnostic* approach to MT regression [5, 6, 10, 13, 24] then consists in leveraging a parameterized task kernel of the form

$$k_{\mathcal{T}}(i, i') = [K_{\text{task}}(b)]_{ii'}, \quad \text{with} \quad K_{\text{task}}(b) = \frac{1}{1+b} I_N + \frac{b}{1+b} \frac{\mathbb{1}_N \mathbb{1}_N^{\top}}{N} \in \mathbb{R}^{N \times N}. \quad (2)$$

Intuitively, parameter $b \geq 0$ governs how similar the tasks are thought to be. When $b = 0$, we have $K_{\text{task}}(b) = I_N$, such that $k_{\mathcal{T}}(i, i') = \delta_{ii'}$, and the tasks are considered to be independent. When b goes to $+\infty$, we have $K_{\text{task}}(b) = \mathbb{1}_N \mathbb{1}_N^{\top} / N$, and all tasks are considered to be one single common task. Any choice of $b \in (0, +\infty)$ corresponds to a tradeoff between these two regimes. We make this intuition explicit in Proposition 2 (Appendix A.1). Note that all quantities depending on the kernel do by definition depend on b . We use the notation $|b$ to make this dependence explicit when relevant.

Given a history of measurements $\{(i_s, x_s), y_s\}_{s=1}^t$, one may then estimate f^{mt} , or equivalently the $\{f_i\}_{i=1}^N$, by standard kernel Ridge regression using the MT kernel k . One obtains the estimates

$$\mu_t(i, x | b) = \mathbf{k}_t(i, x)^{\top} (K_t + \lambda I_t)^{-1} \mathbf{y}_{1:t}, \quad (3)$$

$$\sigma_t^2(i, x | b) = k((i, x), (i, x)) - \mathbf{k}_t(i, x)^{\top} (K_t + \lambda I_t)^{-1} \mathbf{k}_t(i, x), \quad (4)$$

where $\mathbf{k}_t(i, x) = [k((i_s, x_s), (i, x))]_{s=1}^t$, $K_t = [k((i_s, x_s), (i_{s'}, x_{s'}))]_{s, s'=1}^t$, $\mathbf{y}_{1:t} = [y_s]_{s=1}^t$, and $\lambda > 0$ is some regularization parameter. Functions μ_t and σ_t^2 can be interpreted as the posterior mean and variance of a corresponding Gaussian Process model, see [25, 8]. In the next section, we will utilize μ_t and σ_t^2 to construct high-probability confidence intervals for the multitask function f^{mt} .




Figure 2: Novel multi-task confidence width $\beta_t^{\text{new}}(b)$ (see Theorem 1) visualized for large and small values of b . It improves over the naive width $\beta_t^{\text{naive}}(b)$ by a factor of \sqrt{N} at $b = 0$ and as $b \rightarrow +\infty$. Problem parameters were set to $B = 1$, $\epsilon = 0.4$, $N = 20$, $t = 4$, and $\gamma_t^{\text{mt}}(b) = \gamma_t^{\text{st}} = 0$ for all b .

Information gain. An important quantity when analyzing (multitask) kernel regression is the so-called (*multitask*) *information gain*:

$$\gamma_T^{\text{mt}}(b) = \frac{1}{2} \ln |I_T + \lambda^{-1} K_T|.$$

It can be interpreted as the reduction in uncertainty about f^{mt} after having observed a given set of T datapoints. Similarly to single-task setups [25, 8], we use γ_T^{mt} in the next sections to characterize our confidence intervals and regret bounds. Note that γ_T^{mt} depends on the multitask kernel through K_T , and hence on b . In Section 3, we exploit the properties of our multitask kernel to obtain a sharper control over γ_T^{mt} , which is then fundamental to derive improved regret bounds.

2.2 Improved Confidence Intervals

In this section, we utilize the regression estimates obtained in Equations (3) and (4) to construct high probability confidence intervals around the unknown multitask function f^{mt} . First, we assume that $\|f_i\| \leq B$ for all $i \in [N]$, as it is standard in single-task regression. Moreover, let $f_{\text{avg}} = (1/N) \sum_{i=1}^N f_i$ be the average task function, and define

$$\epsilon = \max_i \|f_i - f_{\text{avg}}\| / B. \quad (5)$$

Note that by definition $\epsilon \in [0, 2]$. The smaller ϵ , the more similar the tasks are, the limit case being that all tasks are equal, attained at $\epsilon = 0$. At the other extreme, when $\epsilon \gg 0$ tasks are highly distant and ought to be learned independently. The deviation ϵ plays a crucial role in the subsequent analysis.

A naive confidence interval. As discussed in [6], it is possible to construct the multitask feature map $\tilde{\psi}$ associated to k . One may then rewrite $f^{\text{mt}}(i, x) = \langle \tilde{f}, \tilde{\psi}(i, x) \rangle$, where \tilde{f} is a transformed version of f^{mt} which satisfies $\|\tilde{f}\| \leq B\sqrt{N(1 + b\epsilon^2)}$, see Appendices A.1 and A.2 for details. MT regression thus boils down to single-task regression, over the modified features $\tilde{\psi}(i, x)$, and with target function \tilde{f} . One can then employ well-known linear regression results to obtain confidence intervals for f^{mt} . Using [1, Theorem 3.11, Remark 3.13] and the definition of $\gamma_t^{\text{mt}}(b)$, with probability $1 - \delta$ we have that for all $t \in \mathbb{N}$, $i \in [N]$, and $x \in \mathcal{X}$ it holds $|\mu_t(i, x | b) - f^{\text{mt}}(i, x)| \leq \beta_t^{\text{naive}}(b) \cdot \sigma_t(i, x | b)$, where

$$\beta_t^{\text{naive}}(b) = B\sqrt{N(1 + b\epsilon^2)} + \lambda^{-1/2} \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}.$$

Note that the above confidence interval was already established in [6, Theorem 1]. As expected, it depends on B , N , b , and in a decreasing fashion with respect to ϵ . However, we argue that the above naive choice can be *extremely conservative*. Indeed, when $b = 0$, MT regression treats tasks independently, see Proposition 2. Hence, a valid confidence width from [2, 1, 8] is $\mathcal{O}(B + \sqrt{\gamma_t^{\text{st}}})$, where γ_t^{st} is the single-task maximum information gain. Instead, noting that $\gamma_t^{\text{mt}}(0) = \mathcal{O}(N\gamma_t^{\text{st}})$, see Proposition 1, the naive choice provides $\beta_t^{\text{naive}}(0) = \sqrt{N} \cdot \mathcal{O}(B + \sqrt{\gamma_t^{\text{st}}})$, which is larger by a factor \sqrt{N} . A similar suboptimality gap of \sqrt{N} can also be proven when b tends to $+\infty$. Motivated by the above observation, we derive a novel confidence width that is less conservative than $\beta_t^{\text{naive}}(b)$ for the whole range of possible kernel parameters b .

Theorem 1 (Multitask confidence intervals). *Let $f^{\text{mt}}: [N] \times \mathcal{X} \rightarrow \mathbb{R}$ such that for all $i \in [N]$, $f_i := f^{\text{mt}}(i, \cdot)$ belongs to the RKHS associated to $k_{\mathcal{X}}$ and $\|f_i\| \leq B$. Moreover, let μ_t and σ_t be the*

regression estimates of Equations (3) and (4) with task kernel $k_{\mathcal{T}}(i, j) = [K_{\text{task}}(b)]_{ij}$, parameter $\lambda \in [1/(1+b), 1]$, and noise $\{\xi_{\tau}\}_{\tau=1}^t$ i.i.d. 1-sub-Gaussian. Then, with probability at least $1 - 2\delta$,

$$|\mu_t(i, x | b) - f^{\text{mt}}(i, x)| \leq \beta_t^{\text{new}}(b) \cdot \sigma_t(i, x | b), \quad \forall t \in \mathbb{N}, i \in [N], x \in \mathcal{X},$$

where $\beta_t^{\text{new}}(b) = \min \left\{ \beta_t^{\text{naive}}(b), \beta_t^{\text{small-}b}(b), \beta_t^{\text{large-}b}(b) \right\}$,

$$\beta_t^{\text{small-}b}(b) = B(1 + b\epsilon) \sqrt{\frac{1 + bN}{1 + b}} + \lambda^{-1/2} \sqrt{2(1 + bN)(\gamma_t^{\text{st}} + \ln(N/\delta))},$$

$$\beta_t^{\text{large-}b}(b) = B \sqrt{\frac{(1 + b\epsilon)^2}{1 + b} + \frac{2bN}{1 + b} + \frac{2b(1 + b\epsilon)^2}{N\lambda^2(1 + b)^3} t^2} + \lambda^{-1/2} \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}.$$

The obtained improved confidence width $\beta_t^{\text{new}}(b)$ is the minimum between three confidence widths, see Figure 2. The first one is the naive one $\beta_t^{\text{naive}}(b)$, obtained by standard arguments as outlined above, while $\beta_t^{\text{small-}b}(b)$ and $\beta_t^{\text{large-}b}(b)$ (dashed and dotted lines in Figure 2) are novel and useful for small and large values of b , respectively. Indeed, note that we have $\beta_t^{\text{small-}b}(b) \xrightarrow{b \rightarrow 0} \mathcal{O}(B + \sqrt{\gamma_t^{\text{st}}})$, which is the expected single-task confidence width and \sqrt{N} smaller than $\beta_t^{\text{naive}}(0)$. Similarly, as b goes to $+\infty$ we have $\beta_t^{\text{large-}b}(b) \xrightarrow{b \rightarrow +\infty} \mathcal{O}(B\sqrt{b\epsilon^2 + 2N + 2\epsilon^2 t^2/N}) \xrightarrow{b \rightarrow +\infty} \mathcal{O}(\epsilon B\sqrt{b})$, while $\beta_t^{\text{naive}}(b) \xrightarrow{b \rightarrow +\infty} \mathcal{O}(\epsilon B\sqrt{Nb})$. The obtained confidence width is therefore always smaller than the naive one, but also tighter by a factor \sqrt{N} for the extreme choices $b = 0$ and $b = +\infty$.

From a technical viewpoint, $\beta_t^{\text{small-}b}$ and $\beta_t^{\text{large-}b}$ are obtained by viewing MT regression as a single-task regression over the inflated features $\tilde{\psi}(i, x)$, as also done in [6]. However, unlike [6], we explicitly leverage the expressions of $\tilde{\psi}(i, x | b)$ and $K_{\text{task}}(b)$ as functions of b . In particular, because of the structure of K_{task} , the regression kernel matrix is a rank-one perturbation of a block diagonal matrix, a fact that we exploit, e.g., via Lemma 2. Moreover, we note that refined widths can be obtained if one has access to task-specific constants B_i and ϵ_i . For simplicity of exposition, we focus on uniform (over tasks) B and ϵ . Also, a tighter data-dependent $\beta_t^{\text{large-}b}$ can be utilized as outlined in Appendix A.2.2. Finally, we remark that the obtained multitask intervals do not require i.i.d. data and thus apply to the *adaptive design* setting where data are, e.g., sequentially acquired by the learner, as shown in the next section.

3 New Guarantees for Multitask Online Learning

In this section, we show how the improved confidence interval established in Theorem 1 can be used to derive sharp regret guarantees for multitask online learning. To do so, we also prove novel bounds for the multitask information gain $\gamma_T^{\text{mt}}(b)$. For $t = 1, 2, \dots$ the learning protocol is as follows: nature reveals task index $i_t \in [N]$; the learner chooses strategy $x_t \in \mathcal{X}$ and pays $f^{\text{mt}}(i_t, x_t)$; the learner observes the noisy feedback $y_t = f^{\text{mt}}(i_t, x_t) + \xi_t$. The goal is to minimize for any horizon T the multitask regret

$$R^{\text{mt}}(T) = \sum_{t=1}^T \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) - \sum_{t=1}^T f^{\text{mt}}(i_t, x_t). \quad (6)$$

In the next subsection, we provide a generic algorithm to minimize (6). In particular, we show that naive choices of parameters allow to recover previous approaches with their guarantees, while using the refined confidence width $\beta_t^{\text{new}}(b)$ derived in Theorem 1 yields significant improvements.

3.1 Algorithm and regret guarantees

In line with the online learning literature, our approach is based on the multitask Upper Confidence Bound, defined for any $t \in \mathbb{N}$ as

$$\text{ucb}_t(i, x | b) = \mu_t(i, x | b) + \beta_t(b) \cdot \sigma_t(i, x | b). \quad (7)$$

Here $\beta_t: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a function which assigns a confidence width $\beta_t(b)$ to each kernel parameter b . We consider the general strategy MT-UCB (see Algorithm 1) which, at each round t selects

Algorithm 1 MT-UCB

Require: Domain \mathcal{X} , kernel $k_{\mathcal{X}}$, number of tasks N , width functions $\{\beta_t\}_{t \in \mathbb{N}}$, parameters b, λ .

- 1: **for** $t = 1, 2, \dots$ **do**
- 2: Observe i_t ,
- 3: Play $x_t = \arg \max_{x \in \mathcal{X}} \text{ucb}_{t-1}(i_t, x | b)$,
- 4: Observe $y_t = f^{\text{mt}}(i_t, x_t) + \epsilon_t$,
- 5: Update $\text{ucb}_t(\cdot, \cdot | b)$ based on (3),(4), and (7).

Algorithm	β_t	b	λ
IGP-UCB [8]	$\beta_t^{\text{small-}b}$	0	1
GoB.Lin [6]	β_t^{naive}	b	1
This work	β_t^{new}	b	$\frac{N+b}{N+bN}$

Table 1: Recovering previous works by appropriate choices of β_t, b , and λ .

$x_t = \arg \max_{x \in \mathcal{X}} \text{ucb}_{t-1}(i_t, x | b)$. As summarized in Table 1, both the strategy that runs N independent instances of IGP-UCB (one for each task), and GoB.Lin from [6] are particular cases of MT-UCB. Importantly, whenever $\beta_t(b)$ is set such that $[\mu_t(\cdot, \cdot | b) \pm \beta_t(b) \cdot \sigma_t(\cdot, \cdot | b)]$ is a valid confidence interval for $f^{\text{mt}}(\cdot, \cdot)$, the regret of MT-UCB can be controlled through the following lemma.

Lemma 1. *Suppose that $\lambda \geq (N+b)/(N+bN)$, and that for all tasks i , point x , and time t , we have $f^{\text{mt}}(i, x | b) \in [\mu_t(i, x | b) \pm \beta_t(b) \cdot \sigma_t(i, x | b)]$. Then, the multitask regret of MT-UCB satisfies*

$$R^{\text{mt}}(T) \leq 4 \beta_T(b) \sqrt{\lambda T \gamma_T^{\text{mt}}(b)}.$$

The main novelty of Lemma 1 is that the right-hand side scales with $\lambda^{1/2}$, which might be chosen smaller than 1. This improvement is due to the fact that multitask posterior variances are smaller than $(N+b)/(N+bN) \leq 1$. The right-hand side also depends on the multitask information gain $\gamma_T^{\text{mt}}(b)$, which is nontrivial to compute or upper bound. In the next proposition, we provide practical upper bounds of $\gamma_T^{\text{mt}}(b)$, in terms of the kernel parameter b and the single-task information gain γ_T^{st} .

Proposition 1. *Let $\lambda \leq 1$, $N \geq 2$, and $T_i \geq 1$ for all $i \in [N]$. Then, for any $b \geq 0$, we have*

$$\gamma_T^{\text{mt}}(b) \leq N \gamma_T^{\text{st}} - \frac{Nb}{8(1+b)}, \quad \text{and} \quad \gamma_T^{\text{mt}}(b) \leq \gamma_T^{\text{st}} + \frac{T}{2\lambda(1+b)}.$$

We can now combine Theorem 1, Lemma 1, and Proposition 1 to obtain our main result: a bound on the multitask regret of MT-UCB run with the confidence width β_t^{new} from Theorem 1 and a specific λ .

Theorem 2. *Assume that $B \geq 1$, and that MT-UCB is run with $\beta_t = \beta_t^{\text{new}}$ from Theorem 1, and $\lambda = (N+b)/(N+bN)$. Let $b = N/\epsilon^2$ if $T \leq N$, $b = 1/\epsilon^2$ if $T \geq N$ and $\epsilon \leq N^{-1/4}T^{-1/2}$, and $b = 0$ otherwise. Let $R^{\text{st}}(T) = B\sqrt{T\gamma_T^{\text{st}}} + \sqrt{T\gamma_T^{\text{st}}(\gamma_T^{\text{st}} + \ln(1/\delta))}$ be the single task regret bound achieved by IGP-UCB (up to constant factors). Then, there exists a universal constant C such that with probability $1 - 2\delta$ we have (up to log N factors)*

$$R^{\text{mt}}(T) \leq C \min \left\{ \sqrt{N} R^{\text{st}}(T), R^{\text{st}}(T) + \epsilon B T^{3/2} \left(\sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} + \epsilon \sqrt{T} \right), \right. \\ \left. R^{\text{st}}(T) + \epsilon B T \sqrt{N} \left(\sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} + \epsilon \sqrt{NT} \right) \right\}.$$

The regret bound of Theorem 2 is the minimum between three bounds, obtained exploiting the three different regimes of the confidence width β_t^{new} derived in Theorem 1 (see Figure 2). The **first bound** is obtained using $\beta_t^{\text{new}} \leq \beta_t^{\text{small-}b}$, and shows that our approach cannot be worse than independent learning. Indeed, it can be checked that, when facing N tasks, the regret of running N independent instances of IGP-UCB can be bounded by \sqrt{N} times the single-task regret bound of IGP-UCB, that we denoted by $R^{\text{st}}(T)$. Note however that our analysis slightly differs, insofar as we leverage the multitask information gain, while the independent analysis uses Jensen's inequality to aggregate the individual bounds, see Appendix B for details. Note finally that we are able to recover this bound as $\beta_t^{\text{small-}b}$ is tight at $b = 0$, unlike β_t^{naive} . The **second bound** uses $\beta_t^{\text{new}} \leq \beta_t^{\text{large-}b}$ and consists of two terms: the single task regret bound and an additional term that scales with the task deviation ϵ . When the latter is small, i.e., when tasks are similar, the dominant term is $R^{\text{st}}(T)$, as if only one task were solved. The **third bound** is similar, but obtained using $\beta_t^{\text{new}} \leq \beta_t^{\text{naive}}$ and is useful when $T \geq N$. In

contrast with the independent bound, which does not exploit the task structure, the last two bounds show that multitask learning is always beneficial when the horizon T (and thus the additional ϵ -related term) is small. As expected, this is particularly true when the number of tasks N is large: while the independent bound increases, the second bound *does not depend on N* . On the other hand, one can note that the condition on ϵ to improve over independent becomes more constraining as the horizon T increases. This suggests that the benefit of multitask may vanish with the number of available points per task, an observation which is well-known by practitioners, see e.g. [21]. As far as we know, this work is the first one to provide theoretical evidence of such a phenomenon in online MT learning.

We conclude this section by comparing Theorem 2 to existing results. As already mentioned in the above discussion, independent IGP-UCB is a particular case of MT-UCB, such that we cannot be worse than the independent approach. We incidentally recover its regret bound as the first bound in the minimum of Theorem 2. Regarding GoB.Lin, since it is also a specific instance of MT-UCB (for $\beta_t = \beta_t^{\text{naive}}$ and $\lambda = 1$), Lemma 1 allows to recover its regret bound [6, Theorem 1].

Corollary 1 (Regret of GoB.Lin [6]). *For any b , the multitask regret of GoB.Lin using parameter b satisfies with probability $1 - \delta$*

$$R^{\text{mt}}(T) \leq 4\beta_T^{\text{naive}}(b) \sqrt{T\gamma_T^{\text{mt}}(b)} \leq 6 \left(B\sqrt{N(1+b\epsilon^2)} + \sqrt{\gamma_T^{\text{mt}}(b) + \ln(1/\delta)} \right) \sqrt{T\gamma_T^{\text{mt}}(b)}. \quad (8)$$

If tasks are similar, i.e., when $\epsilon \ll 1$, bound (8) suggests to choose $b > 0$; this does not impact too much the first term, but makes $\gamma_T^{\text{mt}}(b)$ smaller. However, we recall that the above bound instantiated with $b = 0$ does not recover the independent bound. It is instead \sqrt{N} bigger, since β_t^{naive} is not tight at $b = 0$. Hence, the GoB.Lin analysis is not sufficient to show that multitask learning improves over independent learning. Our refined analysis, which uses instead β_t^{new} , closes this gap.

3.2 Adapting to unknown task similarity

In this section, we consider the case where parameter ϵ (i.e., a bound on the task deviation from the average, see (5)) is a-priori unknown. Despite this challenge, we show that the regret bound of Theorem 2 can be approximately attained using an adaptive procedure, AdaMT-UCB (Algorithm 3), relegated to Appendix B.4 due to space limitations. The proposed approach is inspired by the model selection scheme of [22, Section 7] with a few important modifications that we will outline at the end of this section. AdaMT-UCB considers a plausible set of parameters $\mathcal{E} = \{e_1, \dots, e_{|\mathcal{E}|}\} \subset (0, 2]$ and, for each $e \in \mathcal{E}$, initializes an instance of the MT-UCB algorithm with parameters set according to Theorem 2 assuming $\epsilon = e$. We denote such an instance as MT-UCB(e). Moreover, we use the notation ucb_t^e to denote the upper confidence bounds constructed by MT-UCB(e). We assume the existence of some $e \in \mathcal{E}$ such that $e \geq \epsilon$, so that at least one of the learners is *well-specified* (i.e., its confidence bounds contain f^{mt} with high probability). Our goal is to incur a regret which grows as the regret of the learner with the smallest e such that $e \geq \epsilon$, since the smaller the e the smallest the regret bound (see Theorem 2), as long as e is a valid upper bound for ϵ . Let us identify with e^* such learner.

At each round t , AdaMT-UCB uses learner $e_t = \min \mathcal{E}$, and plays the action x_t suggested by it, i.e., the maximizer of $\text{ucb}_t^{e_t}(i_t, \cdot)$. Then, all MT-UCB(e) learners are updated based on the observed reward. In the meantime, a *misspecification test* is carried out to check whether learner e_t is well-specified. Such a test compares the obtained cumulative reward, a lower confidence estimate on such reward according to the other learners, and the believed regret of learner e_t . As long as the test does not trigger, the regret of learner e_t is controlled by the believed one. Instead, if the test triggers, learner e_t can be considered misspecified with high probability. As a result, it gets removed from \mathcal{E} and a *new epoch* starts with the new set \mathcal{E} . Let $\bar{R}_*^{\text{mt}}(T)$ denote the regret bound (Theorem 2) of learner e^* had it been chosen from round 0. We can state the following.

Theorem 3. *Assume that there exists $e \in \mathcal{E}$ such that $e \geq \epsilon$, and let M be the number of learners $e \in \mathcal{E}$ such that $e < \epsilon$ (i.e., the number of misspecified learners in \mathcal{E}). The regret of AdaMT-UCB satisfies with high probability $R^{\text{mt}}(T) = \mathcal{O}(\sqrt{M+1} \cdot \bar{R}_*^{\text{mt}}(T))$.*

Clearly, the number M of misspecified learners is not known in advance but is always less than $|\mathcal{E}|$. Note that when $\epsilon = 0$, we have $M = 0$ and we recover the single task regret bound. Moreover, given $\rho \leq 1$, we show in Appendix B.4 that one can attain a multiplicative accuracy ρ over ϵ , assuming that $\epsilon \geq \epsilon_{\min} > 0$, through an exponential grid with M being polylogarithmic in $1/\rho$ and $1/\epsilon_{\min}$.

Relation with the approach of [22]. Compared to [22, Section 7]—where the goal is to adapt to an unknown features’ dimension—the set of learners considered in AdaMT-UCB share *the same dimension* d . This allows us to exploit the following two novelties with respect to [22]: (1) *all* learners are updated from the data gathered from learner i_t (Line 6 in Algorithm 3), and (2) the lower confidence bounds L^e in the misspecification test (Line 8) are all computed using action x_t (i.e., the action recommended by learner i_t), as opposed to using the actions recommended by each learner e . Both these points are only applicable to our setting, leading to a simpler regret analysis.

4 Multitask Active Learning

The goal of the online learning setup of Section 3 is to optimize the tasks sequentially revealed by nature. In some situations (e.g., in [20] or the drug discovery problem considered in Section 5), however, we care about the performance of multiple tasks *simultaneously*, to eventually learn the best strategy for each one of them. Moreover, we ought to do so with minimal interactions T , i.e., minimizing the queries of the function f^{mt} . We capture this by the following *active learning* protocol.

Learning protocol and regret. At each round t , the learner: chooses a strategy $\{x_t^i, i \in [N]\}$ for each task, chooses which task $i_t \in [N]$ to query, and observes the noisy feedback $y_t = f^{\text{mt}}(i_t, x_t) + \xi_t$. The learner’s goal is to minimize the *active learning* regret:

$$R_{\text{AL}}^{\text{mt}}(T) = \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N \max_{x \in \mathcal{X}} f^{\text{mt}}(i, x) - \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N f^{\text{mt}}(i, x_t^i).$$

Compared to the online learning regret of Equation (6), the learner’s performance at each round is here measured by the average reward coming from *each* task (as opposed to just the task presented by nature). Moreover, compared to online learning, the learner faces the additional challenge of choosing—at each round—from which task information should be gathered. Intuitively, more difficult (or informative) tasks should be queried more often to ensure $R_{\text{AL}}^{\text{mt}}(T)$ grows sublinearly. To the best of our knowledge, the above protocol and regret notion are novel in the multitask literature.

In Algorithm 2 we present MT-AL, an efficient strategy that ensures sublinear active learning regret. Like in MT-UCB, MT-AL constructs confidence intervals around f^{mt} and, at each round, select strategy $x_t^i = \arg \max_{x \in \mathcal{X}} \text{ucb}_{t-1}(i, x)$ for each task $i \in [N]$. When it comes to selecting which task to query, MT-AL selects $i_t \in \arg \max_{i \in [N]} \beta_{t-1}^i \sigma_{t-1}(i, x_t^i)$, i.e., the task for which the believed optimizer x_t^i is subject to maximal uncertainty (we use generic task-dependent widths β_t^i for completeness). This rule, also known as *uncertainty sampling* in the literature [23], intuitively makes sure the learner can control the regrets for the tasks not queried and leads to the following theorem.

Theorem 4. *Suppose that for all tasks i , point x , and time t , we have that $f^{\text{mt}}(i, x) \in [\mu_t(i, x) \pm \beta_t^i \cdot \sigma_t(i, x)]$. Then, the MT-AL algorithm ensures the active learning regret is bounded by*

$$R_{\text{AL}}^{\text{mt}} \leq 2 \sum_{t=1}^T \beta_t^{i_t} \sigma_t(i_t, x_t^{i_t}),$$

where $\{i_t\}$ is the sequence of queried tasks and $\{x_t^{i_t}\}$ the strategies selected for each of them.

The above bound only relies on MT-AL utilizing valid intervals around f^{mt} and thus applies more broadly than our agnostic MT regression, e.g., when such intervals are constructed using a known multitask kernel $k((i, x), (i', x'))$. However, Theorem 4 shows the active learning regret heavily depends on the constructed intervals, similar to online learning. In MT-AL, these are additionally utilized for deciding which task to query at each round. When specialized to our agnostic MT kernel and improved confidence, we obtain the following.

Corollary 2. *Let MT-AL utilize the MT regression estimates of Eq. (3)-(4) with parameters set according to Theorem 2. Moreover, let $\overline{R}^{\text{mt}}(T)$ be the bound on the online learning regret obtained in Theorem 2. Then, with high probability, we have $R_{\text{AL}}^{\text{mt}}(T) \leq \overline{R}^{\text{mt}}(T)$.*




Figure 3: Online and active learning regrets on synthetic and drug discovery MHC-I data, respectively. When utilizing the improved confidence intervals, MT-UCB and MT-AL outperform the other baselines.

Thus, MT-AL ensures the active learning regret is always bounded by its online learning counterpart. Moreover, the same considerations as in Theorem 2 apply also here, regarding the benefit of multitask learning over independent single-task regression for instance.

5 Experiments

The goal of our experiments is to evaluate the effectiveness of the studied MT regression, and in particular of the improved confidence intervals obtained in Section 2, both in online learning and active learning setups. We utilize the following synthetic and real-world data².

Synthetic data: We generate tasks of the form $f_i = (1 - \delta) \cdot \bar{f} + \delta \cdot f_{\text{dev}}^i$, $i \in [N]$, where \bar{f} , f_{dev}^i are random unit vectors representing a common model and individual deviations, respectively. Moreover, actions consist of 10^4 vectors $x \in \mathbb{R}^d$ from the sphere of radius 10. Observation noise is unit normal.

Drug discovery MHC-I data [27]: The goal is to discover the peptides with maximal binding affinity to each Major Histocompatibility Complex class-I (MHC-I) allele. The dataset from [27] contains the standardized binding affinities (IC_{50} values) of different peptide candidates to the MHC-I alleles (tasks). For each allele, the dataset contains ~ 1000 peptides represented as $x \in \mathbb{R}^{45}$ feature vectors. For our experiments, we utilize the 5 alleles A- $\{0201, 0202, 0203, 2301, 2402\}$, since they were shown in [27] to share binding similarity. Note that such a problem falls into our multitask active learning setup, since we would like to retrieve the best peptide for each allele minimizing the number of interactions (i.e., lab experiments). Nevertheless, we also consider its online learning analog where we care about finding the best peptides for each revealed allele.

Online learning. At each round t , a random task $i_t \in [N]$ is observed and point x_t is selected according to the following baselines: (1) *Independent*, which runs N independent IGP-UCB [8] algorithms (corresponding to MT-UCB with $b = 0$), (2) *Single*, which treats all tasks to be the same and runs a unique single-task IGP-UCB (corresponding to MT-UCB with $b = +\infty$), (3) MT-UCB which utilizes an appropriate parameter $0 < b < \infty$ as well as a bound on the tasks similarity ϵ (for synthetic data this can be exactly computed, while for MHC-I data we use $\epsilon = 0.3$) and utilizes the *naive* (i.e., *GoB.Lin*) or *improved* confidence bounds, and (4) AdaMT-UCB which is run with the same b but uses the set of plausible deviations $\mathcal{E} = \{.1, .2, \dots, 1\}$ instead of knowing the true ϵ . For choosing b , we sweep over possible values and select the best-performing one, keeping it fixed for all the baselines.

Active learning. We follow the multitask active learning setup of Section 4. All baselines utilize confidence intervals from the agnostic MT regression of Section 2, where ϵ and b are chosen as for online learning. Moreover, they all utilize the improved confidence intervals, unless otherwise specified. We compare: (1) *Unif.* which chooses the task i_t to be queried uniformly at random (but still selects $x_t^i \in \arg \max_x \text{ucb}_t^i(i, x)$) and employs the *naive* or the *improved* confidence intervals, the offline contextual Bayesian optimization baselines (2) MTS [7] and (3) AE-LSVI [18], and (4) MT-AL which utilizes the *naive* or the *improved* confidence intervals.

²code available at: <https://github.com/sessap/multitask-noregret>.

We report the cumulative regret (online and active learning, respectively) of the considered baselines in Figure 3, averaged over 5 runs. For the synthetic data, we report results for $d = 4$, $N = 5$, $\delta = 0.4$, but provide a full set of experiments for different parameters in Appendix D. In Appendix D we also report the frequencies of each task being queried in our active learning experiments. In Figure 3 (a), both MT-UCB and AdaMT-UCB lead to superior performance compared to the *Independent* and *Single* baselines, demonstrating the benefits of MT regression. In addition, the improved confidence intervals significantly outperform the naive ones. Moreover, we observe AdaMT-UCB achieves comparable (sometimes even better, see Appendix D) performance to MT-UCB. Indeed, instead of using a conservative choice of ϵ , the misspecification test (Line 8 of Algorithm 3) of AdaMT-UCB allows to use a smaller ϵ and only increase it when there is evidence that the constructed intervals do not contain the true tasks. In active learning (Figure 3 (b)), we observe MT-AL has a significant advantage over the uniform sampling baselines and MTS, while performing comparably to AE-LSVI (both methods are similar as discussed in Appendix C.3). Moreover, its regret is bounded by the online learning regret of MT-UCB, conforming with Theorem 4. Importantly, the improved confidence intervals play a crucial role also here and enable a drastic performance improvement compared to the naive ones.

6 Future Directions

We believe this paper opens up several future research directions. The derived confidence intervals, as well as our analysis of the multitask information gain, heavily exploit the structure of the task Gram matrix $K_{\text{task}}(b)$, see Equation (2). However, it remains unclear whether these can be extended to more general kernels. According to the graph perspective of [13], $K_{\text{task}}(b)$ can be seen as $K_{\text{task}}(b) = I_N + L(b)$, where $L(b) \in \mathbb{R}^{N \times N}$ is the Laplacian matrix of a *clique* graph with vertices $[N]$ and edge weight b . Hence, it would be interesting to extend our results to different graph structures. Furthermore, we believe the proposed multitask confidence intervals hold potential for various related domains., e.g., to assess uncertainty in safety-critical systems [3], or to balance exploration-exploitation in multitask reinforcement learning [26], spam filtering [15], or personalized health [14]. In such applications, the introduced notion of active learning regret can serve as a measure of the overall sample efficiency.

Acknowledgements

This work was partially supported by ELSA (European Lighthouse on Secure and Safe AI) funded by the European Union under grant agreement No. 101070617. PL and NCB gratefully acknowledge the financial support from the MUR PRIN grant 2022EKNE5K (Learning in Markets and Society), funded by the NextGenerationEU program within the PNRR scheme (M4C2, investment 1.1), the FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, investment 1.3, line on Artificial Intelligence), and the EU Horizon CL4-2022-HUMAN-02 research and innovation action under grant agreement 101120237, project ELIAS (European Lighthouse of AI for Sustainability).

References

- [1] Abbasi-Yadkori, Y. (2012). Online learning for linearly parametrized control problems. *PhD thesis, Edmonton, Alberta, Canada*.
- [2] Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems (NeurIPS)*, 24.
- [3] Camilleri, R., Wagenmaker, A., Morgenstern, J. H., Jain, L., and Jamieson, K. G. (2022). Active learning with safety constraints. *Advances in Neural Information Processing Systems (NeurIPS)*.
- [4] Caruana, R. (1997). Multitask learning. *Machine learning*, 28:41–75.
- [5] Cavallanti, G., Cesa-Bianchi, N., and Gentile, C. (2010). Linear algorithms for online multitask classification. *Journal of Machine Learning Research*, 11:2901–2934.
- [6] Cesa-Bianchi, N., Gentile, C., and Zappella, G. (2013). A gang of bandits. *Advances in Neural Information Processing Systems (NeurIPS)*, 26.

- [7] Char, I., Chung, Y., Neiswanger, W., Kandasamy, K., Nelson, A. O., Boyer, M., Kolemen, E., and Schneider, J. (2019). Offline contextual bayesian optimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 32.
- [8] Chowdhury, S. R. and Gopalan, A. (2017). On kernelized multi-armed bandits. In *International Conference on Machine Learning (ICML)*, page 844–853.
- [9] Chowdhury, S. R. and Gopalan, A. (2021). No-regret algorithms for multi-task bayesian optimization. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1873–1881.
- [10] Ciliberto, C., Rosasco, L., and Villa, S. (2015). Learning multiple visual tasks while discovering their structure. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 131–139.
- [11] Collobert, R. and Weston, J. (2008). A unified architecture for natural language processing: Deep neural networks with multitask learning. In *International Conference on Machine Learning (ICML)*, pages 160–167.
- [12] Dekel, O., Long, P. M., and Singer, Y. (2007). Online learning of multiple tasks with a shared loss. *Journal of Machine Learning Research*, 8(10).
- [13] Evgeniou, T., Micchelli, C. A., Pontil, M., and Shawe-Taylor, J. (2005). Learning multiple tasks with kernel methods. *Journal of Machine Learning Research*, 6(4).
- [14] Harutyunyan, H., Khachatrian, H., Kale, D. C., Ver Steeg, G., and Galstyan, A. (2019). Multitask learning and benchmarking with clinical time series data. *Scientific data*, 6(1):96.
- [15] Jain, G., Sharma, M., and Agarwal, B. (2019). Spam detection in social media using convolutional and long short term memory neural network. *Annals of Mathematics and Artificial Intelligence*, 85(1):21–44.
- [16] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210.
- [17] Krause, A. and Ong, C. (2011). Contextual Gaussian process bandit optimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 24.
- [18] Li, X., Mehta, V., Kirschner, J., Char, I., Neiswanger, W., Schneider, J., Krause, A., and Bogunovic, I. (2023). Near-optimal policy identification in active reinforcement learning. *International Conference on Learning Representations (ICLR)*.
- [19] Liu, S., Johns, E., and Davison, A. J. (2019). End-to-end multi-task learning with attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1871–1880.
- [20] Lugosi, G., Papaspiliopoulos, O., and Stoltz, G. (2009). Online multi-task learning with hard constraints. In *Conference on Learning Theory (COLT)*.
- [21] Maurer, A., Pontil, M., and Romera-Paredes, B. (2016). The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32.
- [22] Pacchiano, A., Dann, C., Gentile, C., and Bartlett, P. L. (2020). Regret bound balancing and elimination for model selection in bandits and rl. *ArXiv*, abs/2012.13045.
- [23] Settles, B. (2009). Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences.
- [24] Sheldon, D. (2008). Graphical multi-task learning. Technical report, Cornell University, Ithaca, NY.
- [25] Srinivas, N., Krause, A., Kakade, S., and Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. In *International Conference on Machine Learning (ICML)*.

- [26] Vithayathil Varghese, N. and Mahmoud, Q. H. (2020). A survey of multi-task deep reinforcement learning. *Electronics*, 9(9):1363.
- [27] Widmer, C., Toussaint, N. C., Altun, Y., and Rätsch, G. (2010). Inferring latent task structure for multitask learning by multiple kernel learning. *BMC bioinformatics*, 11(8):1–8.

Supplementary Material

We gather here the technical proofs and additional results complementing the main paper.

A Multitask Regression and Improved Confidence Intervals

Notation. In order to coincide with the notation of [6], we introduce the matrix $A(b) \in \mathbb{R}^{N \times N}$ such that $A^{-1}(b) := K_{\text{task}}(b)$. For all $b \geq 0$ we have

$$\begin{aligned} A(b) &= (1+b)I_N - b \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N}, & A(b)^{1/2} &= \sqrt{1+b} I_N + (1 - \sqrt{1+b}) \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N}, \\ A(b)^{-1} &= \frac{1}{1+b} I_N + \frac{b}{1+b} \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N}, & A(b)^{-1/2} &= \frac{1}{\sqrt{1+b}} I_N + \left(1 - \frac{1}{\sqrt{1+b}}\right) \frac{\mathbb{1}_N \mathbb{1}_N^\top}{N}. \end{aligned}$$

Note that in the following we often drop the dependence in b and use only $A, A^{-1}, A^{1/2}, A^{-1/2}$.

A.1 Further intuitions about the MT regression of Section 2.1

Given the history of measurements $\{(i_s, x_s), y_s\}_{s=1}^t$, the agnostic MT regression considered in Section 2.1 estimates f^{mt} (or equivalently the f_i) by standard kernel Ridge regression, i.e., by computing

$$\mu_t = \arg \min_{h \in \mathcal{H}_k} \sum_{s=1}^t (h(i_s, x_s) - y_s)^2 + \lambda \|h\|_{\mathcal{H}_k}^2, \quad (9)$$

where \mathcal{H}_k is the RKHS associated to kernel k , and $\lambda > 0$ some regularization parameter. In the next proposition, we provide some intuition about the role of parameter b by exhibiting regression problems that are equivalent to problem (9) when $b = 0$ and $b = +\infty$. We relegate its proof to the end of this section.

Proposition 2. *Solving the multitask kernel Ridge regression problem (9) for $b = 0$ is equivalent to solve N independent kernel Ridge regressions, using each task's data separately, i.e., we have*

$$\forall i \in [N], \quad \mu_t(i, \cdot | b = 0) = \arg \min_{h \in \mathcal{H}_{k_{\mathcal{X}}}} \sum_{s: i_s = i} (h(x_s) - y_s)^2 + \lambda \|h\|_{\mathcal{H}_{k_{\mathcal{X}}}}^2.$$

On the other hand, solving (9) for $b = +\infty$ is equivalent to solve a unique kernel Ridge regression for all the tasks based on the entire dataset, i.e., we have

$$\forall i \in [N], \quad \mu_t(i, \cdot | b = +\infty) = \arg \min_{h \in \mathcal{H}_{k_{\mathcal{X}}}} \sum_{s=1}^t (h(x_s) - y_s)^2 + \lambda N \|h\|_{\mathcal{H}_{k_{\mathcal{X}}}}^2.$$

Hence, choosing some b in $(0, +\infty)$ is choosing some tradeoff between these two extreme regimes. In Section 3.1, see Theorem 2, we discuss how to set b with respect to the tasks at hand. Proposition 2 provides another intuition: to ensure constant regularization, parameter λ should decrease with b . This key observation was overlooked in [6], partly explaining our better guarantees, see Section 3.1.

Another way to gain intuition about the setting is to compute the feature map induced by the multitask kernel k . Let $\phi: \mathcal{X} \rightarrow \mathbb{R}^d$ be the canonical feature map associated to $k_{\mathcal{X}}$ ³, and $\psi: \mathcal{X} \times [N] \rightarrow \mathbb{R}^{Nd}$ such that $\psi(i, x) = (0, \dots, \phi(x), \dots, 0)$, with non-zero entry at block i . Further define $A_{\otimes} = A \otimes I_d \in \mathbb{R}^{Nd \times Nd}$, where \otimes denote the kronecker product. It is immediate to check that

$$\tilde{\psi}(i, x) = A_{\otimes}^{-1/2} \psi(i, x) = \left(A_{1i}^{-1/2} \phi(x), \dots, A_{ii}^{-1/2} \phi(x), \dots, A_{Ni}^{-1/2} \phi(x) \right) \in \mathbb{R}^{Nd} \quad (10)$$

is a feature map associated to k . Indeed, for any tasks $i, i' \in [N]$, and points $x, x' \in \mathcal{X}$, we have that $\langle \tilde{\psi}(i, x), \tilde{\psi}(i', x') \rangle = \sum_k A_{ik}^{-1/2} A_{i'k}^{-1/2} \langle \phi(x), \phi(x') \rangle = A_{ii'}^{-1} k_{\mathcal{X}}(x, x') = k((i, x), (i', x'))$. Hence, $\tilde{\psi}(i, x)$ stores the feature map $\phi(x)$ in each block j , weighted by some coefficient $A_{ji}^{-1/2}(b)$,

³For the sake of the exposition we consider finite dimensional feature maps. Our results naturally extend to infinite dimensional ones.

which quantifies how similar tasks i and j are assumed to be. When $b = 0$, only block i receives $\phi(x)$. When $b = +\infty$, each block receives $\phi(x)/N$. Denoting abusively $f = (f_1, \dots, f_N)$, note that

$$f^{\text{mt}}(i, x) = \langle f_i, \phi(x) \rangle = \langle f, \psi(i, x) \rangle = \langle A_{\otimes}^{1/2} f, A_{\otimes}^{-1/2} \psi(i, x) \rangle = \langle \tilde{f}, \tilde{\psi}(i, x) \rangle, \quad (11)$$

where $\tilde{f} = (\sum_j A_{1j}^{1/2} f_j, \dots, \sum_j A_{Nj}^{1/2} f_j) \in \mathcal{H}_k$ is a transformed representation of f^{mt} that will play a key role in the subsequent analysis.

Comparison to the model of [9]. We note that an online multitask kernel regression setting has also been considered in [9]. However, the setting of [9] considers *multioutput* functions $f^{\text{mo}}: \mathcal{X} \rightarrow \mathbb{R}^N$ such that the i^{th} output of f^{mo} is given by f_i . The authors learn f^{mo} by leveraging matrix-valued kernels, a multiple output extension of scalar-valued kernel methods. Although the functions modeled by kernel k defined in (1) are isomorphic to that associated to the matrix-valued decomposable kernel given by $k_{\text{MV}}(x, x') = k_{\mathcal{X}}(x, x') A^{-1}$, where we recall that $A_{ii}^{-1} = k_{\mathcal{T}}(i, i')$, we highlight that both models are drastically different. Indeed, while we only observe a single measurement at each time step t , namely a noisy version of $f_{i_t}(x_t)$, the model in [9] assumes that the learner can access N different measurements, that of the $f_i(x_t)$ for all $i \in [N]$. This key difference makes our model more flexible. In particular, in [9] all tasks must be observed the same number of times, and in addition at the same observation points.

A.1.1 Proof of Proposition 2

Recall first that problem (9) writes

$$\mu_t = \arg \min_{h \in \mathcal{H}_k} \sum_{s=1}^t (h(i_s, x_s) - y_s)^2 + \lambda \|h\|_{\mathcal{H}_k}^2.$$

Using identity (11), we obtain

$$\mu_t = \arg \min_{h \in \mathcal{H}_k} \sum_{s=1}^t (\langle h, \tilde{\psi}(i_s, x_s) \rangle - y_s)^2 + \lambda \|h\|_{\mathcal{H}_k}^2,$$

such that standard results for kernel Ridge regression gives that $\mu_t(x) = \langle w_t, \tilde{\psi}(x) \rangle$, with

$$w_t = \left(\sum_{s=1}^t \tilde{\psi}(i_s, x_s) \tilde{\psi}(i_s, x_s)^\top + \lambda I_{Nd} \right)^{-1} \left(\sum_{s=1}^t y_s \tilde{\psi}(i_s, x_s) \right). \quad (12)$$

Now, recall that $\tilde{\psi}(i, x) = A_{\otimes}^{-1/2} \psi(i, x) = (A_{1i}^{-1/2} \phi(x), \dots, A_{ii}^{-1/2} \phi(x), \dots, A_{Ni}^{-1/2} \phi(x))$, such that

$$\tilde{\psi}(i, x | b = 0) = (0, \dots, 0, \underbrace{\phi(x)}_{\text{block } i}, 0, \dots, 0), \quad (13)$$

$$\tilde{\psi}(i, x | b = +\infty) = (\phi(x)/N, \dots, \phi(x)/N). \quad (14)$$

Substituting (13) in (12), we obtain

$$w_t = \begin{pmatrix} \left(\sum_{\substack{s \leq t \\ \text{s.t. } i_s=1}} \phi(x_s) \phi(x_s)^\top + \lambda I_d \right) & & & (0) \\ & \ddots & & \\ & & \left(\sum_{\substack{s \leq t \\ \text{s.t. } i_s=N}} \phi(x_s) \phi(x_s)^\top + \lambda I_d \right) & \\ (0) & & & \end{pmatrix}^{-1} \begin{pmatrix} \sum_{\substack{s \leq t \\ \text{s.t. } i_s=1}} y_s \phi(x_s) \\ \vdots \\ \sum_{\substack{s \leq t \\ \text{s.t. } i_s=N}} y_s \phi(x_s) \end{pmatrix},$$

or again,

$$\forall i \in [N], \quad w_t^{(i)} = \left(\sum_{s \leq t: i_s=i} \phi(x_s) \phi(x_s)^\top + \lambda I_d \right)^{-1} \left(\sum_{s \leq t: i_s=i} y_s \phi(x_s) \right),$$

where $w^{(i)} \in \mathbb{R}^d$ denotes the i^{th} block of concatenated vector $w \in \mathbb{R}^{Nd}$. We thus recover the solutions to the independent regressions stated in the first claim of Proposition 2.

Alternatively, substituting (14) into (12), we obtain

$$w_t = \begin{pmatrix} \frac{1}{N} \sum_{s=1}^t \phi(x_s)\phi(x_s)^\top + \lambda N I_d & & \frac{1}{N} \sum_{s=1}^t \phi(x_s)\phi(x_s)^\top \\ & \ddots & \\ \frac{1}{N} \sum_{s=1}^t \phi(x_s)\phi(x_s)^\top & & \frac{1}{N} \sum_{s=1}^t \phi(x_s)\phi(x_s)^\top + \lambda N I_d \end{pmatrix}^{-1} \begin{pmatrix} \sum_{s=1}^t y_s \phi(x_s) \\ \vdots \\ \sum_{s=1}^t y_s \phi(x_s) \end{pmatrix}$$

$$= \begin{pmatrix} \sum_{s=1}^t \phi(x_s)\phi(x_s)^\top + \lambda N I_d & & (0) \\ & \ddots & \\ (0) & & \sum_{s=1}^t \phi(x_s)\phi(x_s)^\top + \lambda N I_d \end{pmatrix}^{-1} \begin{pmatrix} \sum_{s=1}^t y_s \phi(x_s) \\ \vdots \\ \sum_{s=1}^t y_s \phi(x_s) \end{pmatrix},$$

such that for any $i \in [N]$ we have $w_t^{(i)} = \left(\sum_{s=1}^t \phi(x_s)\phi(x_s)^\top + \lambda N I_d \right)^{-1} \left(\sum_{s=1}^t y_s \phi(x_s) \right)$, which are exactly the solutions to the regression problem on the full dataset stated in the second claim of Proposition 2. Note that the above equality can be easily checked by left multiplying both expressions by $\mathbb{1}_N \mathbb{1}_N^\top \otimes \frac{1}{N} \sum_{s=1}^t \phi(x_s)\phi(x_s)^\top + \lambda N I_{Nd}$. \square

A.2 Proof of Theorem 1

Theorem 1 (Multitask confidence intervals). *Let $f^{\text{mt}}: [N] \times \mathcal{X} \rightarrow \mathbb{R}$ such that for all $i \in [N]$, $f_i := f^{\text{mt}}(i, \cdot)$ belongs to the RKHS associated to $k_{\mathcal{X}}$ and $\|f_i\| \leq B$. Moreover, let μ_t and σ_t be the regression estimates of Equations (3) and (4) with task kernel $k_{\mathcal{T}}(i, j) = [K_{\text{task}}(b)]_{ij}$, parameter $\lambda \in [1/(1+b), 1]$, and noise $\{\xi_\tau\}_{\tau=1}^t$ i.i.d. 1-sub-Gaussian. Then, with probability at least $1 - 2\delta$,*

$$|\mu_t(i, x | b) - f^{\text{mt}}(i, x)| \leq \beta_t^{\text{new}}(b) \cdot \sigma_t(i, x | b), \quad \forall t \in \mathbb{N}, i \in [N], x \in \mathcal{X},$$

where $\beta_t^{\text{new}}(b) = \min \left\{ \beta_t^{\text{naive}}(b), \beta_t^{\text{small-}b}(b), \beta_t^{\text{large-}b}(b) \right\}$,

$$\beta_t^{\text{small-}b}(b) = B(1 + b\epsilon) \sqrt{\frac{1 + bN}{1 + b}} + \lambda^{-1/2} \sqrt{2(1 + bN)(\gamma_t^{\text{st}} + \ln(N/\delta))},$$

$$\beta_t^{\text{large-}b}(b) = B \sqrt{\frac{(1 + b\epsilon)^2}{1 + b} + \frac{2bN}{1 + b} + \frac{2b(1 + b\epsilon)^2}{N\lambda^2(1 + b)^3} t^2} + \lambda^{-1/2} \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}.$$

Proof. Recall that we are interested in obtaining high probability error bounds of the form

$$|\mu_t(i, x) - f^{\text{mt}}(i, x)| \leq \beta_t(b) \cdot \sigma_t(i, x | b)$$

for a suitable choice of confidence width $\beta_t(b)$.

A.2.1 A naive confidence width

Using identity (11) and a direct application of [1, Theorem 3.11 and Remark 3.13], together with the definition of $\gamma_t^{\text{mt}}(b)$, we can select

$$\beta_t = \|\tilde{f}\| + \lambda^{-1/2} \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}.$$

We now upper bound $\|\tilde{f}\|$ explicitly. Recall that

$$\tilde{f} = \left(\sum_j A_{1j}^{1/2} f_j, \dots, \sum_j A_{Nj}^{1/2} f_j \right),$$

with $A = (1 + b)I_N - (b/N)\mathbb{1}\mathbb{1}^\top$, so that we have

$$\begin{aligned}\|\tilde{f}\|^2 &= \sum_{i=1}^N \left\| \sum_{j=1}^N A_{ij}^{1/2} f_j \right\|^2 = \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N A_{ij}^{1/2} A_{ik}^{1/2} \langle f_j, f_k \rangle \\ &= \sum_{j,k=1}^N A_{jk} \langle f_j, f_k \rangle = \sum_{i=1}^N \|f_i\|^2 + b \left(\sum_{i=1}^N \|f_i\|^2 - \frac{1}{N} \sum_{j,k=1}^N \langle f_j, f_k \rangle \right).\end{aligned}\quad (15)$$

On the other side, it holds

$$\begin{aligned}\sum_{i=1}^N \left\| f_i - \frac{1}{N} \sum_{j=1}^N f_j \right\|^2 &= \sum_{i=1}^N \left(\|f_i\|^2 + \frac{1}{N^2} \left\| \sum_{j=1}^N f_j \right\|^2 - \frac{2}{N} \sum_{j=1}^N \langle f_i, f_j \rangle \right) \\ &= \sum_{i=1}^N \|f_i\|^2 + \frac{1}{N} \sum_{i,j=1}^N \langle f_i, f_j \rangle - \frac{2}{N} \sum_{i,j=1}^N \langle f_i, f_j \rangle \\ &= \sum_{i=1}^N \|f_i\|^2 - \frac{1}{N} \sum_{i,j=1}^N \langle f_i, f_j \rangle.\end{aligned}\quad (16)$$

Substituting (16) into (15), we obtain

$$\|\tilde{f}\| = \sum_{i=1}^N \|f_i\|^2 + b \sum_{i=1}^N \|f_i - f_{\text{avg}}\|^2 \leq NB^2(1 + b\epsilon^2).$$

Overall, we can thus choose

$$\beta_t^{\text{naive}}(b) = B\sqrt{N(1 + b\epsilon^2)} + \lambda^{-1/2} \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}. \quad (17)$$

Note that this choice corresponds also to the confidence intervals utilized in the multitask regression setting of [6]. However, we highlight that width (17) is *too conservative*. Indeed, when setting $b = 0$ our model consists in solving single tasks independently, see Proposition 2, such that the confidence width should be of the order $B + \mathcal{O}(\sqrt{\gamma_t^{\text{st}}})$. Instead, we have

$$\beta_t^{\text{naive}}(0) = \sqrt{N}B + \mathcal{O}(\sqrt{\gamma_t^{\text{mt}}(0)}) = \sqrt{N} \left(B + \mathcal{O}(\sqrt{\gamma_t^{\text{st}}}) \right),$$

which is \sqrt{N} bigger. In the next subsection, we derive an improved confidence width which is tight at $b = 0$.

A.2.2 An improved confidence width

As in [8, 2, Proof of Theorem 2], we start by bounding the prediction error as

$$\begin{aligned}|f^{\text{mt}}(i, x) - \mu_t(i, x)| &= |f^{\text{mt}}(i, x) - \mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} (\bar{\mathbf{y}}_{1:t} + \boldsymbol{\xi}_{1:t})| \\ &\leq \underbrace{|f^{\text{mt}}(i, x) - \mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} \bar{\mathbf{y}}_{1:t}|}_{\text{bias error}} + \underbrace{|\mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} \boldsymbol{\xi}_{1:t}|}_{\text{variance error}},\end{aligned}$$

where $\bar{\mathbf{y}}_{1:t} \in \mathbb{R}^t$ is the vector of noise-free outputs such that $\bar{y}_s = f^{\text{mt}}(i_s, x_s) = f_{i_s}(x_s)$, and $\boldsymbol{\xi}_{1:t} \in \mathbb{R}^t$ is the vector of noises. Below, we derive separate bounds for the bias and variance errors. But before doing so, we depart from [8, Proof of Theorem 2] and obtain an explicit lower bound for the predictive variance for task i at point x .

Lower bounding the predictive variance. First, we introduce some notation. Let $A_\otimes \in \mathbb{R}^{Nd \times Nd}$ such that

$$A_\otimes := A \otimes I_d = (1 + b)I_{Nd} - \frac{b}{N} \mathbb{1}_N \mathbb{1}_N^\top \otimes I_d,$$

Moreover, for every $t \in \mathbb{N}$ and $i \in [N]$, let $t_i = \sum_{s \leq t} \mathbb{I}\{i_s = i\}$. We denote by $\Phi_{t_i} \in \mathbb{R}^{t_i \times d}$ the matrix storing in rows the $\phi(x_s)$ for all s such that $i_s = i$. Similarly, let $\Psi_t, \tilde{\Psi}_t \in \mathbb{R}^{t \times Nd}$ storing in

rows the $\psi(i_s, x_s)$ and $\tilde{\psi}(i_s, x_s)$ respectively, for all $s \leq t$. Recall that $\tilde{\psi}(i, x) = A_{\otimes} \psi(i, x)$, such that $\tilde{\Psi}_t = \Psi_t A_{\otimes}$. Further, define

$$M := \Psi_t^{\top} \Psi_t + \lambda(1+b)I_{Nd} = \begin{bmatrix} M_1 & & 0 \\ & \ddots & \\ 0 & & M_N \end{bmatrix}, \quad \text{with } M_i = \Phi_{t_i}^{\top} \Phi_{t_i} + \lambda(1+b)I_d. \quad (18)$$

By properties of the multitask kernel, see (1), the predictive variance for task i at decision point x can be lower bounded as:

$$\begin{aligned} \sigma_t^2(i, x) &= k((i, x), (i, x)) - \mathbf{k}_t(i, x)^{\top} (K_t + \lambda I_t)^{-1} \mathbf{k}_t(i, x) \\ &= \tilde{\psi}(i, x)^{\top} \left(I_{Nd} - \tilde{\Psi}_t^{\top} (\tilde{\Psi}_t \tilde{\Psi}_t^{\top} + \lambda I_t)^{-1} \tilde{\Psi}_t \right) \tilde{\psi}(i, x) \\ &= \tilde{\psi}(i, x)^{\top} \left(I_{Nd} - (\tilde{\Psi}_t^{\top} \tilde{\Psi}_t + \lambda I_{Nd})^{-1} \tilde{\Psi}_t^{\top} \tilde{\Psi}_t \right) \tilde{\psi}(i, x) \end{aligned} \quad (19)$$

$$\begin{aligned} &= \lambda \tilde{\psi}(i, x)^{\top} (\tilde{\Psi}_t^{\top} \tilde{\Psi}_t + \lambda I_{Nd})^{-1} \tilde{\psi}(i, x) \\ &= \lambda \psi(i, x)^{\top} A_{\otimes}^{-1/2} \left(A_{\otimes}^{-1/2} \Psi_t^{\top} \Psi_t A_{\otimes}^{-1/2} + \lambda A_{\otimes}^{-1/2} A_{\otimes} A_{\otimes}^{-1/2} \right)^{-1} A_{\otimes}^{-1/2} \psi(i, x) \\ &= \lambda \psi(i, x)^{\top} (\Psi_t^{\top} \Psi_t + \lambda A_{\otimes})^{-1} \psi(i, x) \\ &= \lambda \psi(i, x)^{\top} \left(M - \frac{\lambda b}{N} \mathbb{1}_N \mathbb{1}_N^{\top} \otimes I_d \right)^{-1} \psi(i, x) \\ &= \lambda \psi(i, x)^{\top} M^{-1} \psi(i, x) \\ &\quad + \lambda \psi(i, x)^{\top} M^{-1} \left(\mathbb{1}_N \mathbb{1}_N^{\top} \otimes \frac{\lambda b}{N} \left(I_{Nd} - \frac{\lambda b}{N} \sum_{i=1}^N M_i^{-1} \right)^{-1} \right) M^{-1} \psi(i, x) \end{aligned} \quad (20)$$

$$\begin{aligned} &= \lambda \phi(x)^{\top} M_i^{-1} \phi(x) + \lambda \phi(x)^{\top} M_i^{-1} \underbrace{\frac{\lambda b}{N} \left(I_d - \frac{\lambda b}{N} \sum_{i=1}^N M_i^{-1} \right)^{-1}}_{:=X^{-1}} M_i^{-1} \phi(x) \\ &= \lambda \left(\|M_i^{-1/2} \phi(x)\|^2 + \frac{\lambda b}{N} \|X^{-1/2} M_i^{-1} \phi(x)\|^2 \right) \\ &\geq \lambda \|M_i^{-1/2} \phi(x)\|^2 + \frac{b}{(1+b)N} \|\lambda X^{-1} M_i^{-1} \phi(x)\|^2 \end{aligned} \quad (21)$$

$$\geq (1+b) \|\lambda M_i^{-1} \phi(x)\|^2 + \frac{b}{(1+b)N} \|\lambda X^{-1} M_i^{-1} \phi(x)\|^2, \quad (22)$$

where (19) comes from the *push-through* equality, (20) from Lemma 2 applied to $D = M$ and $P = -\frac{\lambda b}{N} I_d$, (21) from the fact that $X \succeq 1/(1+b) I_d$, and (22) from $M_i \succeq \lambda(1+b) I_d$. Indeed, the later can easily be checked from (18), which also implies that $(\lambda b)/N \sum_i M_i^{-1} \preceq b/(1+b) I_d$, such that $X \succeq 1/(1+b) I_d$. We now upper bound the bias error in terms of $\sigma_t^2(i, x)$.

Bounding the bias error. Using similar steps as before, we have

$$\begin{aligned} &|f^{\text{mt}}(i, x) - \mathbf{k}_t(i, x)^{\top} (K_t + \lambda I_t)^{-1} \tilde{\mathbf{y}}_{1:t}| \\ &= |\tilde{\psi}(i, x)^{\top} \tilde{f} - \tilde{\psi}(i, x)^{\top} \tilde{\Psi}_t^{\top} (\tilde{\Psi}_t \tilde{\Psi}_t^{\top} + \lambda I_t)^{-1} \tilde{\Psi}_t \tilde{f}| \\ &= |\tilde{\psi}(i, x)^{\top} \tilde{f} - \tilde{\psi}(i, x)^{\top} (\tilde{\Psi}_t^{\top} \tilde{\Psi}_t + \lambda I_{Nd})^{-1} \tilde{\Psi}_t^{\top} \tilde{\Psi}_t \tilde{f}| \\ &= |\lambda \tilde{\psi}(i, x)^{\top} (\Psi_t^{\top} \Psi_t + \lambda I_{Nd})^{-1} \tilde{f}| \\ &= |\lambda \psi(i, x)^{\top} (\Psi_t^{\top} \Psi_t + \lambda A_{\otimes})^{-1} A_{\otimes}^{1/2} \tilde{f}| \\ &= \left| \lambda \psi(i, x)^{\top} \left[M^{-1} + M^{-1} \left(\mathbb{1}_N \mathbb{1}_N^{\top} \otimes \frac{\lambda b}{N} \left(I_d - \sum_{i=1}^N M_i^{-1} \right)^{-1} \right) M^{-1} \right] A_{\otimes}^{1/2} \tilde{f} \right| \end{aligned}$$

$$\begin{aligned}
&= \left| \left(\lambda M_i^{-1} \phi(x) \right)^\top \left[A_{\otimes}^{1/2} \tilde{f} \right]_{[i]} + \frac{\lambda b}{N} \left(\lambda X^{-1} M_i^{-1} \phi(x) \right)^\top \sum_{l=1}^N \left[M^{-1} A_{\otimes}^{1/2} \tilde{f} \right]_{[l]} \right| \\
&\leq \left\| \lambda M_i^{-1} \phi(x) \right\| \cdot \left\| f_i + b(f_i - f_{\text{avg}}) \right\| + \frac{\lambda b}{N} \left\| \lambda X^{-1} M_i^{-1} \phi(x) \right\| \cdot \left\| \sum_{l=1}^N M_l^{-1} \left[A_{\otimes}^{1/2} \tilde{f} \right]_{[l]} \right\| \\
&\leq \sqrt{1+b} \left\| \lambda M_i^{-1} \phi(x) \right\| \frac{B(1+b\epsilon)}{\sqrt{1+b}} \\
&\quad + \sqrt{\frac{b}{(1+b)N}} \left\| \lambda X^{-1} M_i^{-1} \phi(x) \right\| \cdot \lambda \sqrt{\frac{b(1+b)}{N}} \underbrace{\left\| \sum_{l=1}^N M_l^{-1} \left[A_{\otimes}^{1/2} \tilde{f} \right]_{[l]} \right\|}_{:=\|D\|} \\
&\leq \sqrt{\left((1+b) \left\| \lambda M_i^{-1} \phi(x) \right\|^2 + \frac{b \left\| \lambda X^{-1} M_i^{-1} \phi(x) \right\|^2}{(1+b)N} \right) \left(\frac{B^2(1+b\epsilon)^2}{1+b} + \frac{\lambda^2 b(1+b)}{N} \|D\|^2 \right)} \\
\end{aligned} \tag{23}$$

$$\leq \sigma_t(i, x) \cdot \sqrt{\frac{B^2(1+b\epsilon)^2}{1+b} + \frac{\lambda^2 b(1+b)}{N} \|D\|^2}, \tag{24}$$

where we have used $v_{[i]}$ to denote the block i of a concatenated vector in \mathbb{R}^{Nd} , Cauchy-Schwarz inequality to derive (23), and lower bound (22) to obtain (24). Note that D depends on the products between data matrices M_l and task vectors f_1, \dots, f_N , such that it is unknown in general. Below, we provide two upper bounds for $\|D\|$.

Bound 1 (small b range). We can bound $\|D\|$ using Cauchy-Schwarz. We have

$$\|D\| \leq \sum_{l=1}^N \|M_l^{-1}\|_* \left\| \left[A_{\otimes}^{1/2} \tilde{f} \right]_{[l]} \right\| \leq \frac{1}{\lambda(1+b)} \sum_{l=1}^N \|f_l + b(f_l - f_{\text{avg}})\| \leq \frac{NB(1+b\epsilon)}{\lambda(1+b)},$$

which yields

$$\left| f^{\text{mt}}(i, x) - k_t(i, x)^\top (K_t + \lambda I_t)^{-1} \bar{y}_{1:t} \right| \leq \sigma_t(i, x) \cdot B(1+b\epsilon) \sqrt{\frac{1+bN}{1+b}}. \tag{25}$$

The above bound is useful when b is small. Indeed, when $b = 0$, according to the above we recover the single-task confidence width $B \cdot \sigma_t(i, x)$. However, as b goes to $+\infty$, the bound grows as $\mathcal{O}(\sqrt{Nb\epsilon})$, which is an order of \sqrt{b} faster than the naive one, see (17). We thus provide another upper bound on $\|D\|$, which is tighter for large values of b .

Bound 2 (large b range). Alternatively, we can bound $\|D\|$ leveraging the SVD of the data matrices used to build M_l (recall that $M_l = \lambda(1+b)I_d + \Phi_{t_l}^\top \Phi_{t_l}$). For $l \leq N$, let $\Phi_{t_l}^\top \Phi_{t_l} = \sum_k \sigma_k^{(l)} u_k^{(l)} u_k^{(l)\top}$ be the SVD of the data matrix $\Phi_{t_l}^\top \Phi_{t_l}$. We have

$$M_l^{-1} = \sum_k \frac{1}{\sigma_k^{(l)} + \lambda(1+b)} u_k^{(l)} u_k^{(l)\top} = \frac{1}{\lambda(1+b)} I_d - \sum_k \frac{\sigma_k^{(l)}}{\lambda(1+b)(\sigma_k^{(l)} + \lambda(1+b))} u_k^{(l)} u_k^{(l)\top},$$

so that

$$\begin{aligned}
\|D\| &= \left\| \sum_{l=1}^N M_l^{-1} (f_l + b(f_l - f_{\text{avg}})) \right\| \\
&\leq \frac{1}{\lambda(1+b)} \left\| \sum_{l=1}^N f_l \right\| + \frac{1}{\lambda(1+b)} \left\| \sum_{k,l} \frac{\sigma_k^{(l)}}{\sigma_k^{(l)} + \lambda(1+b)} u_k^{(l)} u_k^{(l)\top} (f_l + b(f_l - f_{\text{avg}})) \right\| \\
&\leq \frac{NB}{\lambda(1+b)} + \frac{1}{\lambda(1+b)} \sum_{k,l} \frac{\sigma_k^{(l)}}{\sigma_k^{(l)} + \lambda(1+b)} \|f_l + b(f_l - f_{\text{avg}})\|
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{NB}{\lambda(1+b)} + \frac{B(1+b\epsilon)}{\lambda(1+b)} \sum_{l=1}^N \text{Tr} \left(K_{t_l} (K_{t_l} + \lambda(1+b)I_{t_l})^{-1} \right) \\
&\leq \frac{NB}{\lambda(1+b)} + \frac{B(1+b\epsilon)}{\lambda(1+b)} \sum_{l=1}^N \text{Tr}(K_{t_l}) \cdot \lambda_{\max} \left((K_{t_l} + \lambda(1+b)I_{t_l})^{-1} \right) \\
&\leq \frac{NB}{\lambda(1+b)} + \frac{B(1+b\epsilon)}{\lambda^2(1+b)^2} \sum_{l=1}^N t_l \\
&\leq \frac{NB}{\lambda(1+b)} + \frac{B(1+b\epsilon)}{\lambda^2(1+b)^2} t.
\end{aligned} \tag{26}$$

We note that in practice the data-dependent bound (26) might be tighter than the latter one when λ is small. However, for simplicity of the exposition we focus on the latter data-independent bound. Substituting it into (24), we obtain

$$\begin{aligned}
&|f^{\text{mt}}(i, x) - \mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} \bar{\mathbf{y}}_{1:t}| \\
&\leq \sigma_t(i, x) \cdot \sqrt{\frac{B^2(1+b\epsilon)^2}{1+b} + 2B^2 \frac{bN}{1+b} + 2t^2 \frac{B^2(1+b\epsilon)^2 b}{\lambda^2(1+b)^3 N}} \\
&= \sigma_t(i, x) \cdot B \sqrt{\frac{(1+b\epsilon)^2}{1+b} + \frac{2bN}{1+b} + \frac{2b(1+b\epsilon)^2}{N\lambda^2(1+b)^3} t^2}.
\end{aligned} \tag{27}$$

When b goes to 0, we recover $B \cdot \sigma_t(i, x)$, as for Bound 1. However, the above bound is more useful when b is large. Indeed, when b goes to $+\infty$, we obtain $B\sqrt{b\epsilon^2 + 2N + 2\epsilon^2 t^2/N\lambda^2} \sigma_t(i, x) = \mathcal{O}(B\sqrt{b\epsilon}) \sigma_t(i, x)$, which improves by a factor \sqrt{N} over the $\mathcal{O}(B\sqrt{bN\epsilon}) \sigma_t(i, x)$ term obtained with the naive bound. Recall that when b goes to $+\infty$, MT regression is equivalent to solve a single averaged task based on the whole dataset, see Proposition 2, such that obtaining a confidence width independent from N is expected. Overall, combining (17), (25) and (27), we obtain that

$$|f^{\text{mt}}(i, x) - \mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} \bar{\mathbf{y}}_{1:t}| \leq \beta_t^{\text{bias}}(b) \cdot \sigma_t(i, x), \tag{28}$$

where $\beta_t^{\text{bias}}(b) = B \min \left\{ \sqrt{N(1+b\epsilon^2)}, (1+b\epsilon) \sqrt{\frac{1+bN}{1+b}}, \sqrt{\frac{(1+b\epsilon)^2}{1+b} + \frac{2bN}{1+b} + \frac{2b(1+b\epsilon)^2}{N\lambda^2(1+b)^3} t^2} \right\}$. We now turn to the variance error.

Bounding the variance error. Using (17), the variance error can be naively bounded by $\beta_t^{\text{naive, var}}(b) \cdot \sigma_t(i, x)$, where $\beta_t^{\text{naive, var}}(b) = \lambda^{-1/2} \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}$. However, note that the above bound is conservative when parameter b is small. Indeed, in the limit of $b = 0$ (tasks are treated independently), we know that such error should only depend on the information gain of task i , i.e., $\gamma_{t_i}^{\text{st}}$. Instead, $\gamma_t^{\text{mt}}(0) = \mathcal{O}(N \gamma_{t_i}^{\text{st}})$, which is N times bigger. Let $P_i := \Phi_{t_i}^\top \Phi_{t_i} + I_d$, and note that $M_i^{-1} \preceq P_i^{-1}$, since $\lambda \geq 1/(1+b)$. Looking at a single task i , with probability $1 - \delta$ we have

$$\begin{aligned}
\|M_i^{-1/2} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]}\|^2 &\leq \|P_i^{-1/2} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]}\|^2 \\
&= \boldsymbol{\xi}_{[t_i]}^\top \Phi_{t_i} P_i^{-1} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]} \\
&= \boldsymbol{\xi}_{[t_i]}^\top \Phi_{t_i} (\Phi_{t_i}^\top \Phi_{t_i} + I_d)^{-1} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]} \\
&= \boldsymbol{\xi}_{[t_i]}^\top (K_{t_i} + I_{t_i})^{-1} K_{t_i} \boldsymbol{\xi}_{[t_i]} \\
&= \boldsymbol{\xi}_{[t_i]}^\top (I_{t_i} + K_{t_i}^{-1})^{-1} \boldsymbol{\xi}_{[t_i]} \\
&\leq 2 \left(\frac{1}{2} \ln |I_{t_i} + K_{t_i}| + \ln(1/\delta) \right)
\end{aligned} \tag{29}$$

$$\begin{aligned}
&\leq 2 \left(\frac{1}{2} \ln |I_{t_i} + \lambda^{-1} K_{t_i}| + \ln(1/\delta) \right) \\
&\leq 2(\gamma_{t_i}^{\text{st}} + \ln(1/\delta)),
\end{aligned} \tag{30}$$

where $\boldsymbol{\xi}_{[t_i]} \in \mathbb{R}^{t_i}$ contains the observation noises related to the time steps when task i was active, and $K_{t_i} \in \mathbb{R}^{t_i \times t_i}$ is the individual Gram matrix based on such observations. Equation (29) is obtained by applying [8, Theorem 1] with $\eta = 0$, while (30) derives from $\lambda \leq 1$. By the union bound, we get that with probability at least $1 - \delta$, we have

$$\sup_{i \leq N} \|P_i^{-1/2} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]}\|^2 \leq 2(\gamma_t^{\text{st}} + \ln(N/\delta)).$$

Then, we obtain

$$\begin{aligned} & |k_t(i, x)^\top (K_t + \lambda I_t)^{-1} \boldsymbol{\xi}_{1:t}| \\ &= \left| \tilde{\psi}(i, x)^\top (\tilde{\Psi}_t^\top \tilde{\Psi}_t + \lambda I_{Nd})^{-1} \tilde{\Psi}_t^\top \boldsymbol{\xi}_{1:t} \right| \\ &= \left| \psi(i, x)^\top (\Psi_t^\top \Psi_t + \lambda A_\otimes)^{-1} \Psi_t^\top \boldsymbol{\xi}_{1:t} \right| \\ &= \left| \phi(x)^\top M_i^{-1} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]} + \frac{\lambda b}{N} (X^{-1} M_i^{-1} \phi(x))^\top \sum_{l=1}^N [M^{-1} \Psi_t^\top \boldsymbol{\xi}_{1:t}]_{[l]} \right| \\ &\leq \sqrt{\lambda} \|M_i^{-1/2} \phi(x)\| \cdot \frac{1}{\sqrt{\lambda}} \|M_i^{-1/2} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]}\| \\ &\quad + \sqrt{\frac{b}{(1+b)N}} \|\lambda X^{-1} M_i^{-1} \phi(x)\| \cdot \sqrt{\frac{b(1+b)}{N}} \left\| \sum_{l=1}^N M_l^{-1} \Phi_{t_l}^\top \boldsymbol{\xi}_{[t_l]} \right\| \\ &\leq \sqrt{\lambda} \|M_i^{-1/2} \phi(x)\| \cdot \frac{1}{\sqrt{\lambda}} \|M_i^{-1/2} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]}\| \\ &\quad + \sqrt{\frac{b}{(1+b)N}} \|\lambda X^{-1} M_i^{-1} \phi(x)\| \cdot \sqrt{b(1+b)N} \sup_{l \leq N} \|M_l^{-1/2}\|_* \cdot \|M_l^{-1/2} \Phi_{t_l}^\top \boldsymbol{\xi}_{[t_l]}\| \\ &\leq \sqrt{\lambda} \|M_i^{-1/2} \phi(x)\| \cdot \frac{1}{\sqrt{\lambda}} \|M_i^{-1/2} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]}\| \\ &\quad + \sqrt{\frac{b}{(1+b)N}} \|\lambda X^{-1} M_i^{-1} \phi(x)\| \cdot \sqrt{\frac{bN}{\lambda}} \sup_{l \leq N} \|M_l^{-1/2} \Phi_{t_l}^\top \boldsymbol{\xi}_{[t_l]}\| \\ &\leq \left(\lambda \|M_i^{-1/2} \phi(x)\|^2 + \frac{b}{(1+b)N} \|\lambda X^{-1} M_i^{-1} \phi(x)\|^2 \right)^{1/2} \\ &\quad \cdot \left(\frac{1}{\lambda} \|P_i^{-1/2} \Phi_{t_i}^\top \boldsymbol{\xi}_{[t_i]}\|^2 + \frac{bN}{\lambda} \sup_{l \leq N} \|P_l^{-1/2} \Phi_{t_l}^\top \boldsymbol{\xi}_{[t_l]}\|^2 \right)^{1/2} \\ &\leq \lambda^{-1/2} \sigma_t(i, x) \sup_{l \leq N} \|P_l^{-1/2} \Phi_{t_l}^\top \boldsymbol{\xi}_{[t_l]}\| \sqrt{1 + bN} \\ &\leq \lambda^{-1/2} \sqrt{2(1 + bN)(\gamma_t^{\text{st}} + \ln(N/\delta))} \sigma_t(i, x), \end{aligned} \tag{31}$$

where (31) comes from lower bound (21). Finally, we can take the minimum of this bound and the naive one. Using the union bound again, with probability at least $1 - 2\delta$, we have

$$|k_t(i, x)^\top (K_t + \lambda I_t)^{-1} \boldsymbol{\xi}_{1:t}| \leq \beta_t^{\text{var}}(b) \cdot \sigma_t(i, x), \tag{32}$$

where $\beta_t^{\text{var}}(b) = \lambda^{-1/2} \min \left\{ \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}, \sqrt{2(1 + bN)(\gamma_t^{\text{st}} + \ln(N/\delta))} \right\}$.

Overall error bound. We can obtain the overall prediction error bound by combining bounds (28) and (32) for the bias and variance errors respectively. Hence, with probability $1 - 2\delta$, we have

$$\begin{aligned} |f^{\text{mt}}(i, x) - \mu_t(i, x)| &= |f^{\text{mt}}(i, x) - \mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} (\bar{\mathbf{y}}_{1:t} + \boldsymbol{\xi}_{1:t})| \\ &\leq \underbrace{|f^{\text{mt}}(i, x) - \mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} \bar{\mathbf{y}}_{1:t}|}_{\text{bias error}} + \underbrace{|\mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} \boldsymbol{\xi}_{1:t}|}_{\text{variance error}} \\ &\leq (\beta_t^{\text{bias}}(b) + \beta_t^{\text{var}}(b)) \cdot \sigma_t(i, x), \end{aligned}$$

with $\beta_t^{\text{bias}}(b) = B \min \left\{ \sqrt{N(1+b\epsilon^2)}, (1+b\epsilon) \sqrt{\frac{1+bN}{1+b}}, \sqrt{\frac{(1+b\epsilon)^2}{1+b} + \frac{2bN}{1+b} + \frac{2b(1+b\epsilon)^2}{N\lambda^2(1+b)^3}} t^2 \right\}$, and $\beta_t^{\text{var}}(b) = \lambda^{-1/2} \min \left\{ \sqrt{2(\gamma_t^{\text{mt}}(b) + \ln(1/\delta))}, \sqrt{2(1+bN)(\gamma_t^{\text{st}} + \ln(N/\delta))} \right\}$. In particular, we obtain Theorem 1 by considering specific combinations among the minimums involved in the definitions of $\beta_t^{\text{bias}}(b)$ and $\beta_t^{\text{var}}(b)$. The resulting $\beta_t^{\text{small-}b}(b)$ and $\beta_t^{\text{large-}b}(b)$ are chosen to be tight when b goes to 0 or $+\infty$ respectively. \square

A.2.3 A Kronecker Sherman-Morrison Lemma

We now provide a lemma which extends the Sherman-Morrison formula to kronecker matrices.

Lemma 2. *Let $D_1, \dots, D_N \in \mathbb{R}^{d \times d}$ be invertible, $D = \text{diag}(D_1, \dots, D_N) \in \mathbb{R}^{Nd \times Nd}$, and $P \in \mathbb{R}^{d \times d}$ that commutes with D_i for all $i \in [N]$. Then we have*

$$(D + \mathbb{1}_N \mathbb{1}_N^\top \otimes P)^{-1} = D^{-1} + D^{-1} (\mathbb{1}_N \mathbb{1}_N^\top \otimes Q) D^{-1},$$

where $Q = -\left(I_d + P(D_1^{-1} + \dots + D_N^{-1})\right)^{-1} P = -P \left(I_d + P(D_1^{-1} + \dots + D_N^{-1})\right)^{-1}$.

Proof. It is immediate to check that $(D + \mathbb{1}_N \mathbb{1}_N^\top \otimes P) (D^{-1} + D^{-1} (\mathbb{1}_N \mathbb{1}_N^\top \otimes Q) D^{-1}) = (D^{-1} + D^{-1} (\mathbb{1}_N \mathbb{1}_N^\top \otimes Q) D^{-1}) (D + \mathbb{1}_N \mathbb{1}_N^\top \otimes P) = I_{Nd}$. \square

B New Guarantees for Multitask Online Learning

Details on the independent regret bound. We analyze the regret of the strategy which runs N independent instances of IGP-UCB [8], one per task. Choosing the single-task confidence width $\beta_t^{\text{st}} = B + \sqrt{2(\gamma_t^{\text{st}} + \ln(N/\delta))}$, one can control the individual task regrets and obtain with probability at least $1 - \delta$:

$$\begin{aligned} R^{\text{mt}}(T) &= \sum_{i=1}^N \sum_{t: i_t=i} \max_{x \in \mathcal{X}} f^{\text{mt}}(i, x) - f^{\text{mt}}(i, x_t) \\ &\leq 4 \sum_{i=1}^N \beta_{T_i}^{\text{st}} \sqrt{T_i \gamma_{T_i}^{\text{st}}} \\ &\leq 6 \left(B \sqrt{NT \gamma_T^{\text{st}}} + \sqrt{NT \gamma_T^{\text{st}} (\gamma_T^{\text{st}} + \ln(N/\delta))} \right), \end{aligned}$$

where we have used that $\gamma_{T_i}^{\text{st}} \leq \gamma_T^{\text{st}}$ and Jensen's inequality. We exactly recover the first bound in Theorem 2.

B.1 Proof of Lemma 1

Lemma 1. *Suppose that $\lambda \geq (N+b)/(N+bN)$, and that for all tasks i , point x , and time t , we have $f^{\text{mt}}(i, x | b) \in [\mu_t(i, x | b) \pm \beta_t(b) \cdot \sigma_t(i, x | b)]$. Then, the multitask regret of MT-UCB satisfies*

$$R^{\text{mt}}(T) \leq 4 \beta_T(b) \sqrt{\lambda T \gamma_T^{\text{mt}}(b)}.$$

Proof. The proof follows from standard arguments, see e.g., [25, Theorem 1] and [8, Theorem 3], reproduced here for completeness. For, $i \in [N]$, let $x_i^* = \arg \max_{x \in \mathcal{X}} f^{\text{mt}}(i, x)$. With probability $1 - \delta$, we have

$$\begin{aligned} &\sum_{t=1}^T \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) - f^{\text{mt}}(i_t, x_t) \\ &\leq \sum_{t=1}^T \text{ucb}_{t-1}(i_t, x_{i_t}^* | b) - \text{ucb}_{t-1}(i_t, x_t | b) + 2\beta_t(b) \cdot \sigma_{t-1}(i_t, x_t | b) \end{aligned}$$

$$\begin{aligned}
&\leq 2 \sum_{t=1}^T \beta_t(b) \cdot \sigma_{t-1}(i_t, x_t | b) \\
&\leq 2\beta_T(b) \sqrt{T \sum_{t=1}^T \sigma_{t-1}^2(i_t, x_t | b)}. \tag{33}
\end{aligned}$$

The main twist is that our regularization parameter λ might be smaller than 1, preventing from a direct adaptation of [8, Lemma 4] to bound $\sum_{t=1}^T \sigma_{t-1}^2(i_t, x_t | b)$. However, this happens not to be a problem since our multitask predictive variance are also smaller than in the single-task case. Indeed, as long as $\lambda \geq (N + b)/(N + bN) = A(b)_{ii}^{-1}$, we have for all i and x

$$\begin{aligned}
\sigma_i^2(i, x) &= k((i, x), (i, x)) - \mathbf{k}_t(i, x)^\top (K_t + \lambda I_t)^{-1} \mathbf{k}_t(i, x) \\
&\leq k((i, x), (i, x)) \\
&= A_{ii}^{-1} k_{\mathcal{X}}(x, x) \\
&\leq \lambda.
\end{aligned}$$

Therefore, we have

$$\sum_{t=1}^T \sigma_{t-1}^2(i_t, x_t) = \lambda \sum_{t=1}^T \lambda^{-1} \sigma_{t-1}^2(i_t, x_t) \leq 2\lambda \sum_{t=1}^T \ln(1 + \lambda^{-1} \sigma_{t-1}^2(i_t, x_t)) \leq 4\lambda \gamma_T^{\text{mt}}(b), \tag{34}$$

where we have used that $x \leq \ln(1 + x)$ for any $x \in [0, 1]$, applied to the $\lambda^{-1} \sigma_{t-1}^2(i_t, x_t) \leq 1$, and [8, Lemma 3]. Substituting (34) into (33) concludes the proof. \square

B.2 Proof of Proposition 1

Proposition 1. *Let $\lambda \leq 1$, $N \geq 2$, and $T_i \geq 1$ for all $i \in [N]$. Then, for any $b \geq 0$, we have*

$$\gamma_T^{\text{mt}}(b) \leq N\gamma_T^{\text{st}} - \frac{Nb}{8(1+b)}, \quad \text{and} \quad \gamma_T^{\text{mt}}(b) \leq \gamma_T^{\text{st}} + \frac{T}{2\lambda(1+b)}.$$

Proof. Recall that the multitask kernel writes $k((i, x), (i', x')) = k_{\mathcal{T}}(i, i') \cdot k_{\mathcal{X}}(x, x')$. Hence, the multitask Gram matrix K_T can be written as $K_T = K_{\mathcal{T}} \odot K_{\mathcal{X}}$, where $K_{\mathcal{T}}, K_{\mathcal{X}} \in \mathbb{R}^{T \times T}$ are the task (respectively domain) Gram matrices. Moreover, up to rearranging the rows and columns of $K_{\mathcal{T}}$ (which does not change the determinant), we can assume that points are ordered by task activations. Let $T_i = \sum_{t=1}^T \mathbb{1}\{i_t = i\}$ be the number of times task i has been queried. We have

$$\begin{aligned}
K_{\mathcal{T}} &= \begin{pmatrix} \frac{b+N}{(1+b)N} \mathbb{1}_{T_1} \mathbb{1}_{T_1}^\top & & \frac{b}{(1+b)N} \mathbb{1}_{T_1} \mathbb{1}_{T-T_1}^\top \\ & \ddots & \\ \frac{b}{(1+b)N} \mathbb{1}_{T_N} \mathbb{1}_{T-T_N}^\top & & \frac{b+N}{(1+b)N} \mathbb{1}_{T_N} \mathbb{1}_{T_N}^\top \end{pmatrix} \\
&= \frac{b}{1+b} \frac{\mathbb{1}_T \mathbb{1}_T^\top}{N} + \frac{1}{1+b} \begin{pmatrix} \mathbb{1}_{T_1} \mathbb{1}_{T_1}^\top & & 0 \\ & \ddots & \\ 0 & & \mathbb{1}_{T_N} \mathbb{1}_{T_N}^\top \end{pmatrix}.
\end{aligned}$$

We also introduce the block notation $K_{\mathcal{X}}^{(i,j)} \in \mathbb{R}^{T_i \times T_j}$ and $K_{\mathcal{X}}^{\text{diag}} \in \mathbb{R}^{T \times T}$ such that

$$K_{\mathcal{X}} = \begin{pmatrix} K_{\mathcal{X}}^{(i,j)} \end{pmatrix}, \quad \text{and} \quad K_{\mathcal{X}}^{\text{diag}} = \begin{pmatrix} K_{\mathcal{X}}^{(1,1)} & & 0 \\ & \ddots & \\ 0 & & K_{\mathcal{X}}^{(N,N)} \end{pmatrix}.$$

Our bounds are based on the observation that $g: t \mapsto \ln |tX + (1-t)Y|$ is concave, with derivative $g'(t) = \text{Tr} \left((tX + (1-t)Y)^{-1} (X - Y) \right)$. Applying the concavity inequality at $t = 0$ and $t = 1$, we obtain that for any positive semi-definite matrices X and Y and any $t \in [0, 1]$ we have

$$\ln |tX + (1-t)Y| \leq \ln |Y| + t \cdot \left(\text{Tr}(Y^{-1}X) - T \right), \tag{35}$$

$$\ln |tX + (1-t)Y| \leq \ln |X| + (1-t) \cdot (\text{Tr}(X^{-1}Y) - T). \quad (36)$$

Hence, for any $b \geq 0$ we have

$$\begin{aligned} 2\gamma_T^{\text{mt}}(b) &= \ln |I_T + \lambda^{-1}K_T| \\ &= \ln |I_T + \lambda^{-1}K_T \odot K_{\mathcal{X}}| \\ &= \ln \left| I_T + \lambda^{-1} \left(\frac{b}{1+b} \frac{K_{\mathcal{X}}}{N} + \frac{1}{1+b} K_{\mathcal{X}}^{\text{diag}} \right) \right| \\ &= \ln \left| \frac{b}{1+b} \left(I_T + \lambda^{-1} \frac{K_{\mathcal{X}}}{N} \right) + \frac{1}{1+b} \left(I_T + \lambda^{-1} K_{\mathcal{X}}^{\text{diag}} \right) \right| \\ &\leq \ln |I_T + \lambda^{-1} K_{\mathcal{X}}^{\text{diag}}| + \frac{b}{1+b} \left[\text{Tr} \left(\left(I_T + \lambda^{-1} K_{\mathcal{X}}^{\text{diag}} \right)^{-1} \left(I_T + \lambda^{-1} \frac{K_{\mathcal{X}}}{N} \right) \right) - T \right] \end{aligned} \quad (37)$$

$$= 2 \sum_{i=1}^N \gamma_{T_i}^{\text{st}} + \frac{b}{1+b} \left[\text{Tr} \left(\left(I_T + \lambda^{-1} K_{\mathcal{X}}^{\text{diag}} \right)^{-1} \left(I_T + \lambda^{-1} \frac{K_{\mathcal{X}}}{N} \right) \right) - T \right], \quad (38)$$

where (37) derives from (35). We now take a closer look at the second term. We have

$$\begin{aligned} &\text{Tr} \left(\left(I_T + \lambda^{-1} K_{\mathcal{X}}^{\text{diag}} \right)^{-1} \left(I_T + \lambda^{-1} \frac{K_{\mathcal{X}}}{N} \right) \right) \\ &= \text{Tr} \left(\begin{pmatrix} \left(I_{T_1} + \lambda^{-1} K_{\mathcal{X}}^{(1,1)} \right)^{-1} & & 0 \\ & \ddots & \\ 0 & & \left(I_{T_N} + \lambda^{-1} K_{\mathcal{X}}^{(N,N)} \right)^{-1} \end{pmatrix} \left(I_T + \lambda^{-1} \frac{K_{\mathcal{X}}}{N} \right) \right) \\ &= \sum_{i=1}^N \text{Tr} \left(\left(I_{T_i} + \lambda^{-1} K_{\mathcal{X}}^{(i,i)} \right)^{-1} \left(I_{T_i} + \lambda^{-1} \frac{K_{\mathcal{X}}^{(i,i)}}{N} \right) \right) \\ &= \sum_{i=1}^N \sum_{\tau=1}^{T_i} \frac{\lambda + \sigma_{\tau}^{(i)}/N}{\lambda + \sigma_{\tau}^{(i)}} \end{aligned} \quad (39)$$

where $\{\sigma_{\tau}^{(i)}\}_{\tau \leq T_i}$ are the eigenvalues of $K_{\mathcal{X}}^{(i,i)}$, possibly equal to 0. For any i , let $F^{(i)}: \mathbb{R}^{T_i} \rightarrow \mathbb{R}$ the functions which to any $\sigma = (\sigma_1, \dots, \sigma_{T_i})$ associates $F^{(i)}(\sigma) = \sum_{\tau=1}^{T_i} (\lambda + \sigma_{\tau}/N)/(\lambda + \sigma_{\tau})$. For any τ_0, τ_1 , we have

$$\begin{aligned} &(\sigma_{\tau_0} - \sigma_{\tau_1}) \left(\frac{\partial F^{(i)}(\sigma)}{\partial \sigma_{\tau_0}} - \frac{\partial F^{(i)}(\sigma)}{\partial \sigma_{\tau_1}} \right) \\ &= (\sigma_{\tau_0} - \sigma_{\tau_1}) \left(\frac{(1/N)(\lambda + \sigma_{\tau_0}) - (\lambda + \sigma_{\tau_0}/N)}{(\lambda + \sigma_{\tau_0})^2} - \frac{(1/N)(\lambda + \sigma_{\tau_1}) - (\lambda + \sigma_{\tau_1}/N)}{(\lambda + \sigma_{\tau_1})^2} \right) \\ &= \lambda \frac{N-1}{N} (\sigma_{\tau_1} - \sigma_{\tau_0}) \left(\frac{1}{(\lambda + \sigma_{\tau_0})^2} - \frac{1}{(\lambda + \sigma_{\tau_1})^2} \right) \\ &\geq 0, \end{aligned}$$

such that $F^{(i)}$ is Schur-convex. Hence, (39) is maximized at eigenvalues of the form $(T_i, 0, \dots, 0)$, since the latter majorizes any other admissible distribution of the eigenvalues (recall that we must have $\sigma_{\tau}^{(i)} \geq 0$ and $\sum_{\tau=1}^{T_i} \sigma_{\tau}^{(i)} = \text{Tr}(K_{\mathcal{X}}^{(i,i)}) \leq T_i$), with value

$$\sum_{i=1}^N (T_i - 1) + \frac{\lambda + T_i/N}{\lambda + T_i} \leq T - N + N \frac{1 + 1/N}{2} \leq T - \frac{N}{4},$$

where we have used that $\lambda \leq 1$, $T_i \geq 1$, and $N \geq 2$. Substituting into (38), we finally obtain

$$\gamma_T^{\text{mt}}(b) \leq N\gamma_T^{\text{st}} - \frac{Nb}{8(1+b)}. \quad (40)$$

The second bound is obtained by modifying (37). Instead, we now apply (36) and obtain

$$\begin{aligned}
2\gamma_T^{\text{mt}}(b) &\leq \ln \left| I_T + \lambda^{-1} \frac{K_{\mathcal{X}}}{N} \right| + \frac{1}{1+b} \left[\text{Tr} \left(\left(I_T + \lambda^{-1} \frac{K_{\mathcal{X}}}{N} \right)^{-1} \left(I_T + \lambda^{-1} K_{\mathcal{X}}^{\text{diag}} \right) \right) - T \right] \\
&\leq 2\gamma_T^{\text{st}} + \frac{1}{1+b} \left(\text{Tr} \left(I_T + \lambda^{-1} K_{\mathcal{X}}^{\text{diag}} \right) - T \right) \\
&\leq 2\gamma_T^{\text{st}} + \frac{T}{\lambda(1+b)}
\end{aligned} \tag{41}$$

where (41) comes from von Neumann's trace inequality and is tight when $N \rightarrow +\infty$. \square

B.3 Proof of Theorem 2

Theorem 2. Assume that $B \geq 1$, and that MT-UCB is run with $\beta_t = \beta_t^{\text{new}}$ from Theorem 1, and $\lambda = (N+b)/(N+bN)$. Let $b = N/\epsilon^2$ if $T \leq N$, $b = 1/\epsilon^2$ if $T \geq N$ and $\epsilon \leq N^{-1/4}T^{-1/2}$, and $b = 0$ otherwise. Let $R^{\text{st}}(T) = B\sqrt{T}\gamma_T^{\text{st}} + \sqrt{T}\gamma_T^{\text{st}}(\gamma_T^{\text{st}} + \ln(1/\delta))$ be the single task regret bound achieved by IGP-UCB (up to constant factors). Then, there exists a universal constant C such that with probability $1 - 2\delta$ we have (up to log N factors)

$$\begin{aligned}
R^{\text{mt}}(T) &\leq C \min \left\{ \sqrt{N}R^{\text{st}}(T), R^{\text{st}}(T) + \epsilon BT^{3/2} \left(\sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} + \epsilon\sqrt{T} \right), \right. \\
&\quad \left. R^{\text{st}}(T) + \epsilon BT\sqrt{N} \left(\sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} + \epsilon\sqrt{NT} \right) \right\}.
\end{aligned}$$

Proof. From Lemma 1 and the choice $\beta_t = \beta_t^{\text{new}}$, we have with probability at least $1 - 2\delta$

$$R^{\text{mt}}(T) \leq 4\beta_T^{\text{new}}(b)\sqrt{\lambda T \gamma_T^{\text{mt}}(b)}.$$

First bound, recovering independent learning. Hence, in particular, we have

$$\begin{aligned}
R^{\text{mt}}(T) &\leq 4\beta_T^{\text{small-}b}(b)\sqrt{\lambda T \gamma_T^{\text{mt}}(b)} \\
&= 4 \left(B(1+b\epsilon)\sqrt{\frac{1+bN}{1+b}} + \lambda^{-1/2}\sqrt{2(1+bN)(\gamma_T^{\text{st}} + \ln(N/\delta))} \right) \sqrt{\lambda T \gamma_T^{\text{mt}}(b)} \\
&\leq 6 \left(B\frac{1+b\epsilon}{1+b}\sqrt{\frac{(1+bN)(b+N)}{N}} + \sqrt{(1+bN)(\gamma_T^{\text{st}} + \ln(N/\delta))} \right) \sqrt{T \left(N\gamma_T^{\text{st}} - \frac{Nb}{8(1+b)} \right)},
\end{aligned}$$

where we have used the first claim of Proposition 1 and the choice $\lambda = (N+b)/(N+bN)$. Substituting $b = 0$ in the above equation, we recover the independent learning bound, i.e.,

$$R^{\text{mt}}(T) \leq 6 \left(B\sqrt{NT\gamma_T^{\text{st}}} + \sqrt{NT\gamma_T^{\text{st}}(\gamma_T^{\text{st}} + \ln(N/\delta))} \right). \tag{42}$$

Hence, even in the least favorable cases, the multitask approach allows to recover the independent baseline by using $b = 0$. Note that this is only made possible by the fact that $\beta_t^{\text{small-}b}$ is tight at $b = 0$.

A first bound for small ϵ . Here, we instead use the bound

$$\begin{aligned}
R^{\text{mt}}(T) &\leq 4\beta_T^{\text{naive}}(b)\sqrt{\lambda T \gamma_T^{\text{mt}}(b)} \\
&= 4 \left(B\sqrt{N(1+b\epsilon^2)} + \lambda^{-1/2}\sqrt{2(\gamma_T^{\text{mt}}(b) + \ln(1/\delta))} \right) \sqrt{\lambda T \gamma_T^{\text{mt}}(b)} \\
&\leq 4 \left(B\sqrt{\frac{(1+b\epsilon^2)(b+N)}{1+b}} + \sqrt{2 \left(\gamma_T^{\text{st}} + \frac{NT}{b} + \ln(1/\delta) \right)} \right) \sqrt{T \left(\gamma_T^{\text{st}} + \frac{NT}{b} \right)},
\end{aligned}$$

where we have used the second claim of Proposition 1 and the choice $\lambda = (N + b)/(N + bN)$. Substituting $b = 1/\epsilon^2$ in the above equation, we obtain

$$\begin{aligned}
\frac{R^{\text{mt}}(T)}{6} &\leq \left(B\sqrt{\frac{N+1/\epsilon^2}{1+1/\epsilon^2}} + \sqrt{\gamma_T^{\text{st}} + \epsilon^2 NT + \ln(1/\delta)} \right) \sqrt{T(\gamma_T^{\text{st}} + \epsilon^2 NT)} \\
&\leq \left(B(1 + \epsilon\sqrt{N}) + \sqrt{\gamma_T^{\text{st}} + \epsilon^2 NT + \ln(1/\delta)} \right) \sqrt{T(\gamma_T^{\text{st}} + \epsilon^2 NT)} \\
&\leq \sqrt{T\gamma_T^{\text{st}}} \left(B + \sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} \right) + \epsilon \left[B\sqrt{NT\gamma_T^{\text{st}}} + T\sqrt{N\gamma_T^{\text{st}}} + B\sqrt{NT} \right. \\
&\quad \left. + \epsilon BNT + \epsilon NT^{3/2} + T\sqrt{N(\gamma_T^{\text{st}} + \ln(1/\delta))} \right] \\
&= \sqrt{T\gamma_T^{\text{st}}} \left(B + \sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} \right) + \mathcal{O} \left(\epsilon BT\sqrt{N(\gamma_T^{\text{st}} + \ln(1/\delta))} + \epsilon^2 BNT^{3/2} \right). \tag{43}
\end{aligned}$$

Bound (43) is composed of two terms: the single-task regret, and an additional term which scales with the deviation ϵ to the average task. Hence, as ϵ goes to 0, i.e., when tasks get more similar, we adaptively recover the single-task bound. Moreover, note that (43) is smaller than (42) as long as $\epsilon \leq 1/(N^{1/4}\sqrt{T})$. Hence, by choosing $b = (1/\epsilon^2) \cdot \mathbb{1}\{\epsilon \leq 1/(N^{1/4}\sqrt{T})\}$, we can obtain the minimum of the two bounds.

A second bound for small ϵ . Finally, we can use that

$$\begin{aligned}
R^{\text{mt}}(T) &\leq 4\beta_T^{\text{large-}b}(b)\sqrt{\lambda T\gamma_T^{\text{mt}}(b)} \\
&= 4 \left(B\sqrt{\frac{(1+b\epsilon)^2}{1+b} + \frac{2bN}{1+b} + \frac{2b(1+b\epsilon)^2}{N\lambda^2(1+b)^3} T^2} + \lambda^{-1/2}\sqrt{2(\gamma_T^{\text{mt}}(b) + \ln(1/\delta))} \right) \sqrt{\lambda T\gamma_T^{\text{mt}}(b)} \\
&\leq 6 \left(B\sqrt{\frac{(1+b\epsilon)^2(b+N)}{(1+b)^2N} + \frac{2b(b+N)}{(1+b)^2} \left(1 + \frac{(1+b\epsilon)^2}{(b+N)^2} T^2 \right)} + \sqrt{\gamma_T^{\text{st}} + \frac{NT}{b} + \ln(1/\delta)} \right) \\
&\quad \cdot \sqrt{T \left(\gamma_T^{\text{st}} + \frac{NT}{b} \right)},
\end{aligned}$$

where we have used the second claim of Proposition 1 and the choice $\lambda = (N + b)/(N + bN)$. Choosing $b = N/\epsilon^2$, we have

$$\begin{aligned}
\frac{(1+b\epsilon)^2(b+N)}{(1+b)^2N} &= \frac{(1+\frac{N}{\epsilon})^2(1+\frac{1}{\epsilon^2})}{(1+\frac{N}{\epsilon^2})^2} \leq \frac{2(1+\frac{N^2}{\epsilon^2})}{(1+\frac{N}{\epsilon^2})} \frac{1+\epsilon^2}{N+\epsilon^2} \leq 2N\frac{5}{N} = 10, \\
\frac{2b(b+N)}{(1+b)^2} &\leq \frac{2(b+N)}{(1+b)} = 2\frac{N+\frac{N}{\epsilon^2}}{1+\frac{N}{\epsilon^2}} = 2N\frac{1+\epsilon^2}{N+\epsilon^2} \leq 10, \\
\frac{1+b\epsilon}{b+N} &= \frac{b\epsilon + N\epsilon - N\epsilon + 1}{b+N} \leq \epsilon + \frac{1}{\frac{N}{\epsilon^2} + N} \leq \epsilon + \frac{\epsilon^2}{N} \leq 3\epsilon.
\end{aligned}$$

Substituting in the above bound, we obtain

$$\begin{aligned}
R^{\text{mt}}(T) &= \mathcal{O} \left(B\sqrt{1 + \epsilon^2 T^2} + \sqrt{\gamma_T^{\text{st}} + \epsilon^2 T + \ln(1/\delta)} \right) \sqrt{T(\gamma_T^{\text{st}} + \epsilon^2 T)} \\
&= \mathcal{O} \left(\sqrt{T\gamma_T^{\text{st}}} \left(B + \sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} \right) + \epsilon BT^{3/2}\sqrt{\gamma_T^{\text{st}}} \right. \\
&\quad \left. + \epsilon T \left(B\sqrt{1 + \epsilon^2 T^2} + \sqrt{\gamma_T^{\text{st}} + \epsilon^2 T + \ln(1/\delta)} \right) \right)
\end{aligned}$$

Algorithm 3 AdaMT-UCB

Require: Finite set $\mathcal{E} \subset (0, 2]$, learning rate $\eta > 0$.

- 1: For $e \in \mathcal{E}$, initialize learner MT-UCB(e) with $b, \lambda, \{\beta_t\}_{t \in \mathbb{N}}$ according to Theorem 2 using $\epsilon = e$.
 - 2: Initialize $\tau = U = R = 0, L^e = 0, \forall e \in \mathcal{E}$.
 - 3: **for** $t=1, \dots$, **do**
 - 4: Choose learner $e_t = \min_{e \in \mathcal{E}}$.
 - 5: Observe i_t and play action from MT-UCB(e_t), i.e. $x_t = \arg \max_{x \in \mathcal{X}} \text{ucb}_{t-1}^{e_t}(i_t, x)$.
 - 6: Observe: $y_t = f^{\text{mt}}(i_t, x_t) + \xi_t$, and update *all* learners $\{\text{MT-UCB}(e)\}_{e \in \mathcal{E}}$ based on observation.
 - 7: Accumulate:

$$\tau += 1, \quad U += y_t, \quad R += 2\beta_{t-1}^{e_t} \sigma_{t-1}^{e_t}(i_t, x_t), \quad L^e += \text{lcb}_{t-1}^e(i_t, x_t), \forall e \in \mathcal{E}.$$
 - 8: Misspecification test:
$$U + R + c\sqrt{\tau \ln(\ln(\tau)/\delta)} < \max_{e \in \mathcal{E}} L^e \quad (46)$$
 - 9: **if** condition (46) is true **then**
 - 10: # Learner e_t is misspecified w.h.p.. Hence, start a new epoch.
 - 11: $\mathcal{E} = \mathcal{E} \setminus \{e_t\}$ and reset $\tau = U = R = 0, L^e = 0, \forall e \in \mathcal{E}$.
-

$$= \mathcal{O} \left(\sqrt{T\gamma_T^{\text{st}}} \left(B + \sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} \right) + \epsilon BT^{3/2} \sqrt{\gamma_T^{\text{st}} + \ln(1/\delta)} + \epsilon^2 BT^2 \right). \quad (44)$$

Again, bound (44) is composed of two terms: the single-task regret, and an additional term which goes to 0 as ϵ goes to 0. Interestingly, this last part is independent from N , which is a consequence of $\beta_t^{\text{large-}b}$ being \sqrt{N} smaller than β_t^{naive} at $b = +\infty$. For small values of T , namely when $T \leq N$, (44) is thus smaller than (43). Choosing $b = N/\epsilon^2$ when $T \leq N$, and as before otherwise, ensures to obtain the minimum of (42), (43), and (44). \square

B.4 Adapting to unknown tasks' similarity

In Algorithm 3 we summarize the AdaMT-UCB approach discussed in Section 4. In the misspecification test (Line 8), lcb_t^e are lower confidence bound functions, defined as:

$$\text{lcb}_t^e(i, x) = \mu_t(i, x | b^e) - \beta_t(b^e) \cdot \sigma_t(i, x | b^e),$$

where b^e is the kernel parameter chosen (according to Theorem 2) by each learner e . Moreover, c is an absolute constant such that, by standard concentration arguments and for all times τ ,

$$\left| U - \sum_{t=1}^{\tau} f^{\text{mt}}(i_t, x_t) \right| \leq c\sqrt{\tau \ln(\ln(\tau)/\delta)}, \quad (45)$$

with probability $1 - \delta$, see, e.g., [22, Lemma B.1].

B.4.1 Proof of Theorem 3

Theorem 3. *Assume that there exists $e \in \mathcal{E}$ such that $e \geq \epsilon$, and let M be the number of learners $e \in \mathcal{E}$ such that $e < \epsilon$ (i.e., the number of misspecified learners in \mathcal{E}). The regret of AdaMT-UCB satisfies with high probability $R^{\text{mt}}(T) = \mathcal{O}(\sqrt{M+1} \cdot \overline{R}_*^{\text{mt}}(T))$.*

Among the set of learners defined by parameters $e \in \mathcal{E}$, we have identified with e^* the (well-specified) learner with the smallest e such that $e \geq \epsilon$. Then, our goal is to obtain a regret bound which grows as the regret of learner e^* . We have denoted with $\overline{R}_*^{\text{mt}}(T)$ the regret bound of such learner had it been chosen from time 0.

We first prove the following auxiliary lemma, which is the analog of [22, Theorem 7.1].

Lemma 3. *With probability at least $1 - \delta$, the misspecification test in Equation (46) does not trigger if all learners in \mathcal{E} are well-specified and their confidence intervals contain f^{mt} .*

Proof. When all learners $e \in \mathcal{E}$ are well-specified and their intervals contain the true function, for all t it holds $\text{lcb}_t^e(i_t, x_t) \leq \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x)$. Thus, for each learner $e \in \mathcal{E}$, with probability at least

$1 - \delta$,

$$\begin{aligned}
L^e &= \sum_{t=1}^T \text{lcb}_{t-1}^e(i_t, x_t) \\
&\leq \sum_{t=1}^T \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) \\
&\leq \sum_{t=1}^T f^{\text{mt}}(i_t, x_t) + 2\beta_{t-1}^{e_t} \sigma_{t-1}^{e_t}(i_t, x_t) \\
&\leq U + c\sqrt{T \ln(\ln(T)/\delta)} + R,
\end{aligned}$$

where, in addition to Equation (45), we have used that $\sum_t \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) - f^{\text{mt}}(i_t, x_t) \leq \sum_t \text{ucb}_t(i_t, x_t) - \text{lcb}_t(i_t, x_t) = \sum_t 2\beta_{t-1}^{e_t} \sigma_{t-1}^{e_t}(i_t, x_t)$. Thus, the misspecification test of (46) does not trigger. \square

Let us now bound the overall regret of AdaMT-UCB. First, we can decompose it into the regrets inside each epoch:

$$\begin{aligned}
R^{\text{mt}}(T) &= \sum_{t=1}^T \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) - f^{\text{mt}}(i_t, x_t) \\
&= \sum_{s=1}^{\# \text{ of Epochs}} \sum_{t \in \text{Epoch-}s} \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) - f^{\text{mt}}(i_t, x_t) \\
&= \sum_{s=1}^{\# \text{ of Epochs}} R_s^{\text{mt}}(T_s),
\end{aligned}$$

where we have defined T_s to be the duration of epoch s and $R_s^{\text{mt}}(T_s)$ its corresponding regret. Note that by Lemma 3, the maximum number of terminated epochs corresponds to the number of misspecified learners in the initial set \mathcal{E} . Thus, letting M be such number, with high probability:

$$R^{\text{mt}}(T) \leq \sum_{s=1}^{M+1} R_s^{\text{mt}}(T_s). \quad (47)$$

During each epoch. Let us now look at what happens during each epoch s . For simplicity, we will condition on the event that the intervals of learner e^* contain the true f^{mt} , and on the event of Equation (45). Note that by definition, during each epoch the misspecification test has not triggered. In particular, this is true when testing against learner e^* . That is,

$$U + R + c\sqrt{T_s \ln(\ln(T_s)/\delta)} \geq \sum_{t \in \text{Epoch-}s} \text{lcb}_{t-1}^{e^*}(i_t, x_t),$$

which, by letting $e_s = \min_{e \in \mathcal{E}} e_s$ be the learner utilized in epoch s , implies:

$$\sum_{t \in \text{Epoch-}s} f^{\text{mt}}(i_t, x_t) + 2\beta_{t-1}^{e_s} \sigma_{t-1}^{e_s}(i_t, x_t) + 2c\sqrt{T_s \ln(\ln(T_s)/\delta)} \geq \sum_{t \in \text{Epoch-}s} \text{lcb}_{t-1}^{e^*}(i_t, x_t). \quad (48)$$

Then, using the above condition:

$$\begin{aligned}
R_s^{\text{mt}}(T_s) &- 2c\sqrt{T_s \ln \frac{\ln(T_s)}{\delta}} \\
&= \sum_{t \in \text{Epoch-}s} \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) - f^{\text{mt}}(i_t, x_t) - 2c\sqrt{T_s \ln \frac{\ln(T_s)}{\delta}} \\
\text{(By Eq. (48))} &\leq \sum_{t \in \text{Epoch-}s} \max_{x \in \mathcal{X}} f^{\text{mt}}(i_t, x) - \text{lcb}_{t-1}^{e^*}(i_t, x_t) + \sum_{t \in \text{Epoch-}s} 2\beta_{t-1}^{e_s} \sigma_{t-1}^{e_s}(i_t, x_t)
\end{aligned}$$

$$\begin{aligned}
(e^* \text{ is well-specified}) &\leq \sum_{t \in \text{Epoch-}s} 2\beta_{t-1}^{e^*} \sigma_{t-1}^{e^*}(i_t, x_t) + \sum_{t \in \text{Epoch-}s} 2\beta_{t-1}^{e_s} \sigma_{t-1}^{e_s}(i_t, x_t) \\
&\leq 4\beta_{T_s}(b^{e^*}) \sqrt{\lambda(b^{e^*}) T_s \gamma_{T_s}^{\text{mt}}(b^{e^*})} + 4\beta_{T_s}(b^{e_s}) \sqrt{\lambda(b^{e_s}) T_s \gamma_{T_s}^{\text{mt}}(b^{e_s})} \\
&\leq 2\overline{R}_*^{\text{mt}}(T_s).
\end{aligned}$$

In the last two inequalities, we have used the same proof steps of Lemma 1 to bound the sum of confidence widths for learners e^* and e_s , and finally, the fact that $e_s \leq e^*$ and thus the regret bound of learner e_s is bounded by $\overline{R}_*^{\text{mt}}(T_s)$ (since the bound from Theorem 2 increases with ϵ).

Overall, combining the latter with Equation (47), we obtain

$$\begin{aligned}
R^{\text{mt}}(T) &\leq \sum_{s=1}^{M+1} 2\overline{R}_*^{\text{mt}}(T_s) + 2c\sqrt{T_s \ln \frac{\ln(T_s)}{\delta}} \\
&= 2 \sum_{s=1}^{M+1} C^*(T_s) \sqrt{T_s} + c\sqrt{T_s \ln \frac{\ln(T_s)}{\delta}} \\
&\leq 2C^*(T) \sqrt{T} \sqrt{M+1} + 2c\sqrt{(M+1)T \ln \frac{\ln(T)}{\delta}} \\
&= 2\overline{R}_*^{\text{mt}}(T) \sqrt{M+1} + 2c\sqrt{(M+1)T \ln \frac{\ln(T)}{\delta}}.
\end{aligned}$$

where we have use the fact that MT-UCB regret bounds are of the form $\overline{R}_*^{\text{mt}}(T_s) = C^*(T_s) \sqrt{T_s}$ for some appropriate function $C^*(T_s)$, see Theorem 2, and Jensen's inequality. \square

How many learners are needed? Let \mathcal{E} be the exponential grid $\{1, \rho, \rho^2, \dots, \rho^{M-1}\}$, for some $\rho < 1$. Let $\epsilon \in [0, 2]$ be the true tasks similarity parameter. By definition, the best learner is better than the learner m^* such that $\rho^{m^*+1} \leq \epsilon \leq \rho^{m^*}$. The estimate it uses for ϵ is better than $\epsilon^* := \rho^{m^*}$, which satisfies $\epsilon/\epsilon^* \in [\rho, 1]$. Hence, the bigger ρ , the more precise we are. However, the number of learners needed also increases with ρ . Indeed, if we want to be able to identify up to ϵ_{\min} , we have to choose M such that

$$\rho^{M-1} \leq \epsilon_{\min} \quad \text{or again} \quad M \geq 1 + \frac{\log(1/\epsilon_{\min})}{\log(1/\rho)}.$$

C Active Learning

C.1 Proof of Theorem 4

Theorem 4. *Suppose that for all tasks i , point x , and time t , we have that $f^{\text{mt}}(i, x) \in [\mu_t(i, x) \pm \beta_t^i \cdot \sigma_t(i, x)]$. Then, the MT-AL algorithm ensures the active learning regret is bounded by*

$$R_{\text{AL}}^{\text{mt}} \leq 2 \sum_{t=1}^T \beta_t^{i_t} \sigma_t(i_t, x_t^{i_t}),$$

where $\{i_t\}$ is the sequence of queried tasks and $\{x_t^{i_t}\}$ the strategies selected for each of them.

Proof. Let $x_*^i \in \arg \max_{x \in \mathcal{X}} f^{\text{mt}}(i, x)$. Then, the active learning regret of MT-AL can be bounded as

$$\begin{aligned}
R_{\text{AL}}^{\text{mt}}(T) &:= \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N f^{\text{mt}}(i, x_*^i) - \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N f^{\text{mt}}(i, x_t^i) \\
&\leq \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N \text{ucb}_{t-1}(i, x_*^i) - \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N \text{ucb}_{t-1}(i, x_t^i) + 2 \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N \beta_{t-1}^i \sigma_{t-1}(i, x_t^i) \\
&\leq 2 \sum_{t=1}^T \frac{1}{N} \sum_{i=1}^N \beta_{t-1}^i \sigma_{t-1}(i, x_t^i)
\end{aligned}$$

$$\leq 2 \sum_{t=1}^T \beta_{t-1}^i \sigma_{t-1}(i_t, x_t^{i_t}) \underbrace{\frac{1}{N} \sum_{i=1}^N 1}_{=1}$$

The first inequality holds since by assumption, for all tasks i , point x , and time t , we have $f^{\text{mt}}(i, x) \in [\mu_t(i, x) \pm \beta_t^i \cdot \sigma_t(i, x)]$. The second one follows since, at each round t MT-AL select $x_t^i = \arg \max_x \text{ucb}_{t-1}(i, x)$ for all i , and the third one since $i_t = \arg \max_i \beta_{t-1}^i \sigma_{t-1}(i, x_t^i)$. \square

C.2 Proof of Corollary 2

Corollary 2. *Let MT-AL utilize the MT regression estimates of Eq. (3)-(4) with parameters set according to Theorem 2. Moreover, let $\overline{R}^{\text{mt}}(T)$ be the bound on the online learning regret obtained in Theorem 2. Then, with high probability, we have $R_{\text{AL}}^{\text{mt}}(T) \leq \overline{R}^{\text{mt}}(T)$.*

Proof. First, since MT-AL utilizes the MT regression estimates of Eq. (3)-(4) with parameters set according to Theorem 2, with high probability $f^{\text{mt}}(i, x) \in [\mu_t(i, x) \pm \beta_t^i \cdot \sigma_t(i, x)]$ and the results of Theorem 4 holds. Then, the result follows since, according to the proof of Theorem 2, for every sequence of revealed tasks $\{i_t\}_{t=1}^T$ (in particular the ones chosen by MT-AL), it holds $2 \sum_{t=1}^T \beta_{t-1}^i \sigma_{t-1}(i_t, x_t^{i_t}) \leq \overline{R}^{\text{mt}}(T)$, see Appendix B. \square

C.3 Comparison with AE-LSVI [18]

The proposed MT-AL algorithm can be compared with the offline contextual Bayesian algorithm AE-LSVI [18] whose goal is to quickly discover the optimal strategy for each context (task, in our case). Like MT-AL, AE-LSVI selects strategy $x_t^i = \arg \max_x \text{ucb}_{t-1}(i, x)$ for each task i . However, unlike MT-AL, AE-LSVI queries the task $i_t = \arg \max_{i \in [N]} [\text{ucb}_t(i, x_t^i) - \max_{x \in \mathcal{X}} \text{lcb}_t(i, x)]$. This, together with a terminal rule for the final reported actions $\{x_T^i\}_{i=1}^N$, ensures the latter are $\mathcal{O}(\beta_T \sqrt{\gamma_T^{\text{mt}}}/\sqrt{T})$ -approximate optimal for each task. We note that a similar error can also be proven when using MT-AL, by turning the active learning regret bound into a last-iterate approximation error (effectively dividing the regret by T). However, unlike our approach, it is unclear whether the active learning regret of AE-LSVI is sublinear. In particular, by querying the task with maximal uncertainty MT-AL controls the regrets for all the other tasks (see last inequality in Proof of Theorem 4). Instead, the query strategy of AE-LSVI considers a truncated uncertainty, since $[\text{ucb}_t(i, x_t^i) - \max_{x \in \mathcal{X}} \text{lcb}_t(i, x)] \leq [\text{ucb}_t(i, x_t^i) - \text{lcb}_t(i, x_t^i)] = \beta_t^i \sigma_t(i, x_t^i)$.

D Additional experimental results

All our experiments were run using 8 CPUs at 3.7 GHz.

In Figure 4 we report online learning an active learning regrets of additional synthetic experiments for different values of N (number of tasks) and parameter δ (inversely proportional to the tasks' similarity, see Section 5). We observe that the improvement of MT regression over independent learning increases with the number of tasks and with their similarity (i.e., for small δ). Moreover, as N increases, the improved confidence intervals further improve over the naive ones, as well as the benefit of active learning compared to uniform sampling. All these considerations conform with our theory.

In our active learning experiments, we further analyze the frequencies of each task being queried by the considered baselines. In Figure 5, we report such frequencies averaged over 10 runs for a random instance of our synthetic experiments with $N = 5$ and $\delta = 0.4$, i.e. the setup of Figure 3 (b). We observe the active learning baselines (MT-AL, MTS, and AE-LSVI) query more frequently task-1 and task-5, which are found to be the tasks with the smallest norm $\|f_i\|$ and thus with the smallest signal-to-noise ratio. In Figure 6, we report the query frequencies for our drug discovery experiments averaged over 5 runs. We observe MT-AL and AE-LSVI (which both query tasks with maximal uncertainty) query slightly more often task-3 (corresponding to allele A-0203). Instead, MTS (which query the task yielding the maximum performance improvement) focuses significantly more often task-1 (allele A-0201). This produces high rewards in the early rounds but incurs linearly growing regret overall (see Figure 3 (b)).

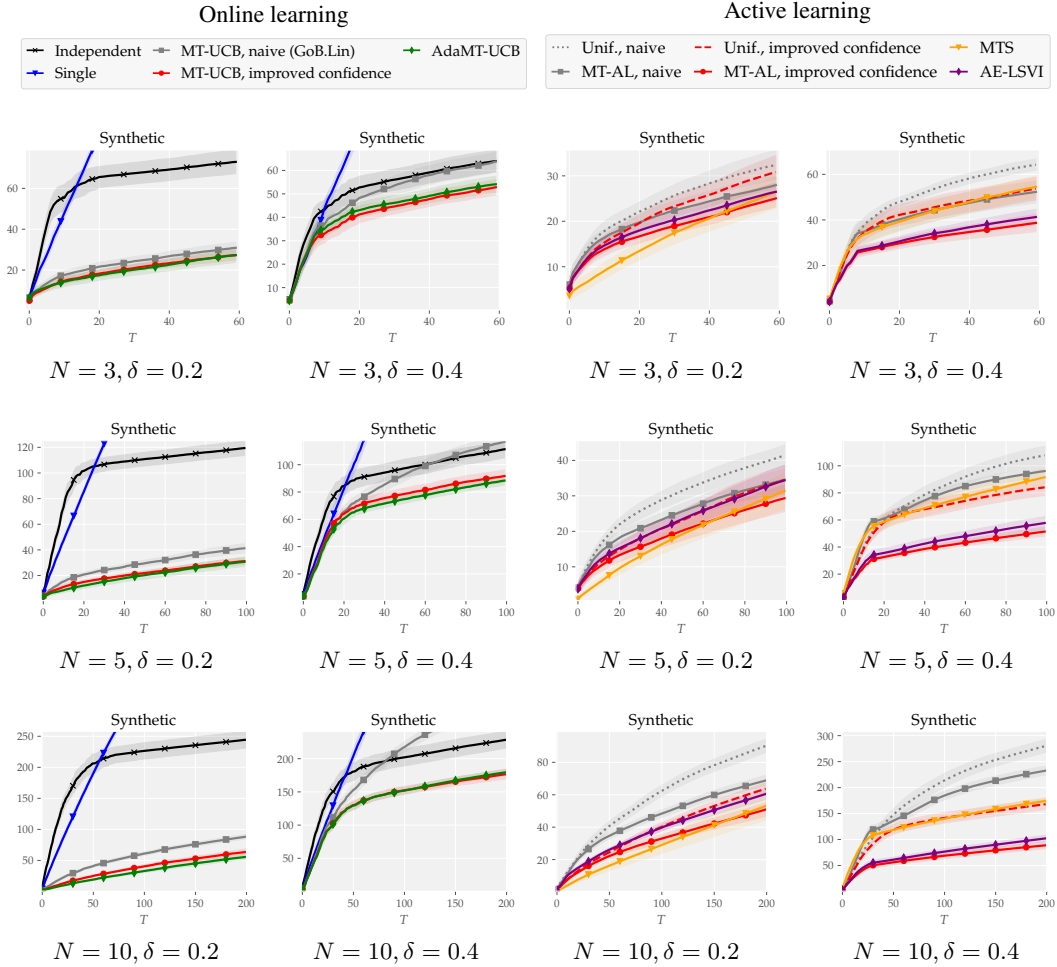


Figure 4: Online (left) and active (right) learning regrets of synthetic experiments for different parameters: N is the number of tasks while parameter δ is inversely proportional to the tasks' similarity, see Section 5 for how task vectors are generated.

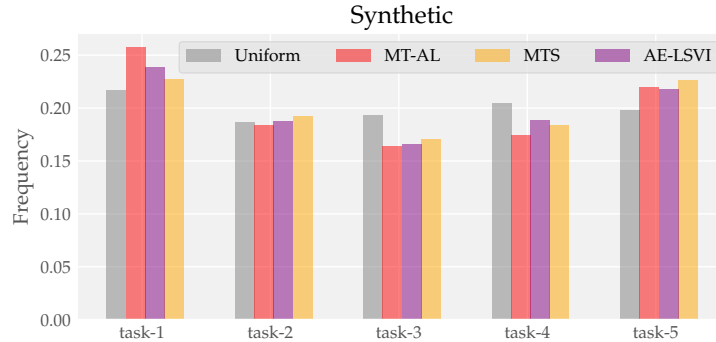


Figure 5: Frequency each task being queried by each of the considered baselines. The plot refers to our synthetic experiments with $N = 5$, $\delta = 0.4$, i.e. the setup of Figure 3 (b). We generate a single set of task vectors $\{f_i\}_{i=1}^N$ and average results over 10 runs, each with horizon $T = 100$.

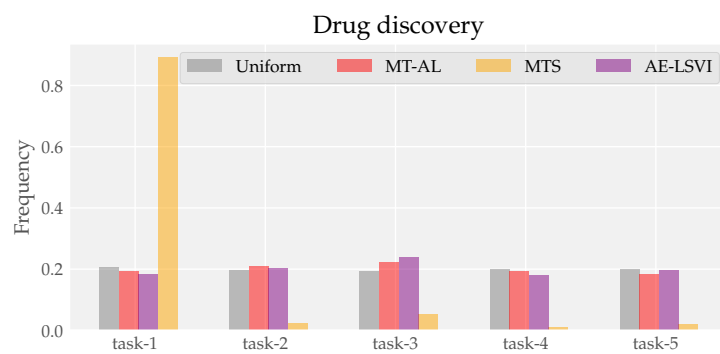


Figure 6: Frequency of each task being queried by each of the considered baselines for our drug discovery experiments (results averaged over 5 runs, each with horizon $T = 600$).