
Analysing Multi-Task Regression via Random Matrix Theory with Application to Time Series Forecasting

Romain Ilbert^{*1,2} Malik Tiomoko¹ Cosme Louart³ Ambroise Odonnat^{1,4}
Vasilii Feofanov¹ Themis Palpanas² Ievgen Redko¹

¹Huawei Noah's Ark Lab, Paris, France ²LIPADE, Paris Descartes University, Paris, France

³School of Data Science, The Chinese University of Hong Kong, Shenzhen, China

⁴Inria, Univ. Rennes 2, CNRS, IRISA

Abstract

In this paper, we introduce a novel theoretical framework for multi-task regression, applying random matrix theory to provide precise performance estimations, under high-dimensional, non-Gaussian data distributions. We formulate a multi-task optimization problem as a regularization technique to enable single-task models to leverage multi-task learning information. We derive a closed-form solution for multi-task optimization in the context of linear models. Our analysis provides valuable insights by linking the multi-task learning performance to various model statistics such as raw data covariances, signal-generating hyperplanes, noise levels, as well as the size and number of datasets. We finally propose a consistent estimation of training and testing errors, thereby offering a robust foundation for hyperparameter optimization in multi-task regression scenarios. Experimental validations on both synthetic and real-world datasets in regression and multivariate time series forecasting demonstrate improvements on univariate models, incorporating our method into the training loss and thus leveraging multivariate information.

1 Introduction

The principle of multi-task learning, which encompasses concepts such as "learning to learn" and "knowledge transfer," offers an efficient paradigm that mirrors human intelligence. Unlike traditional machine learning, which tackles problems on a task-specific basis with tailored algorithms, multi-task learning leverages shared information across various tasks to enhance overall performance. This approach not only improves accuracy but also addresses challenges related to insufficient data and data cleaning. By analyzing diverse and multimodal data and structuring representations multi-task learning can significantly boost general intelligence, similar to how humans integrate skills. This concept is well-established [33] and has been successfully applied to a wide range of domains, such as computer vision [41], natural language processing [35, 37] and biology [17, 28, 42].

Our Method. We consider the multi-task regression framework [5] with T tasks with the input space $\mathcal{X}^{(t)} \subset \mathbb{R}^d$ and the output space $\mathcal{Y}^{(t)} \subset \mathbb{R}^q$ for $t \in \{1, \dots, T\}$. For each task t , we assume that we are given n_t training examples organized into the feature matrix $\mathbf{X}^{(t)} = [\mathbf{x}_1^{(t)}, \dots, \mathbf{x}_{n_t}^{(t)}] \in \mathbb{R}^{d \times n_t}$ and the corresponding response matrix $\mathbf{Y}^{(t)} = [\mathbf{y}_1^{(t)}, \dots, \mathbf{y}_{n_t}^{(t)}] \in \mathbb{R}^{q \times n_t}$, where $\mathbf{x}_i^{(t)} \in \mathcal{X}^{(t)}$ represents the i -th feature vector of the t -th task and $\mathbf{y}_i^{(t)} \in \mathcal{Y}^{(t)}$ is the associated response. In order to have a more tractable and insightful setup, we follow Romera-Paredes et al. [36] and consider the linear multi-task regression model. In particular, we study a straightforward linear signal-plus-noise model

*Correspondence to: romain.ilbert@hotmail.fr.

that evaluates the response $y_i^{(t)}$ for the i -th sample of the t -th task as follows:

$$\mathbf{Y}^{(t)} = \frac{\mathbf{X}^{(t)\top} \mathbf{W}_t}{\sqrt{Td}} + \boldsymbol{\varepsilon}^{(t)}, \quad \forall t \in \{1, \dots, T\}, \quad (1)$$

where $\boldsymbol{\varepsilon}^{(t)} \in \mathbb{R}^{n_t \times q}$ is a matrix of noise vectors with each $\boldsymbol{\varepsilon}_i^{(t)} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_N)$, $\boldsymbol{\Sigma}_N \in \mathbb{R}^{q \times q}$ denoting the covariance matrix. The matrix $\mathbf{W}_t \in \mathbb{R}^{d \times q}$ denotes the signal-generating hyperplane for task t . We denote the concatenation of all task-specific hyperplanes by $\mathbf{W} = [\mathbf{W}_1^\top, \dots, \mathbf{W}_T^\top]^\top \in \mathbb{R}^{Td \times q}$. We assume that \mathbf{W}_t can be decomposed as a sum of a common matrix $\mathbf{W}_0 \in \mathbb{R}^{d \times q}$, which captures the shared information across all the tasks, and a task-specific matrix $\mathbf{V}_t \in \mathbb{R}^{d \times q}$, which captures deviations specific to the task t :

$$\mathbf{W}_t = \mathbf{W}_0 + \mathbf{V}_t. \quad (2)$$

Given the multitask regression framework and the linear signal-plus-noise model, we now want to retrieve the common and specific hyperplanes, \mathbf{W}_0 and \mathbf{V}_t , respectively. To achieve this, we study the following minimization problem governed by a parameter λ that controls the balance between the common and specific components of \mathbf{W} :

$$\mathbf{W}_0^*, \{\mathbf{V}_t^*\}_{t=1}^T, \lambda^* = \operatorname{argmin} \quad \frac{1}{2\lambda} \|\mathbf{W}_0\|_F^2 + \frac{1}{2} \sum_{t=1}^T \frac{\|\mathbf{V}_t\|_F^2}{\gamma_t} + \frac{1}{2} \sum_{t=1}^T \left\| \mathbf{Y}^{(t)} - \frac{\mathbf{X}^{(t)\top} \mathbf{W}_t}{\sqrt{Td}} \right\|_F^2. \quad (3)$$

where $\gamma = [\gamma_1, \dots, \gamma_T]$ is a vector of task-specific regularization hyperparameters. Each γ_t controls how much the model overfits (small γ_t) or underfits (large γ_t) on task t .

Contributions and Main Results. We seek to provide a rigorous theoretical study of (3) in high dimension, providing practical insights that make the theoretical application implementable in practice. Our contributions can be summarized in four major points:

1. We provide exact train and test risks for the solutions of (3) using random matrix theory. We then decompose the test risk into a signal and noise term responsible for the effectiveness of multi-task learning and the negative transfer, respectively.
2. We show how the signal and noise terms compete with each other depending on λ for which we obtain an optimal value optimizing the test risk.
3. In a particular case, we derive a closed-form data-dependent solution for the optimal λ^* , involving raw data covariances, signal-generating hyperplane, noise levels, and the size and number of data sets.
4. We offer empirical estimates of these model variables to develop a ready-to-use method for hyperparameter optimization in multitask regression problems within our framework.

Applications. As an application, we view multivariate time series forecasting (MTSF) as a multi-task problem and show how the proposed regularization method can be used to efficiently employ univariate models in the multivariate case. Firstly, we demonstrate that our method improves PatchTST [32] and DLinear [53] compared to the case when these approaches are applied independently to each variable. Secondly, we show that the univariate models enhanced by the proposed regularization achieve performance similar to the current state-of-the-art multivariate forecasting models such as SAMformer [19] and iTransformer [24].

2 Related Work

Usual Bounds in Multi-Task Learning. Since its introduction [1, 5, 26], multi-task learning (MTL) has demonstrated that transfer learning concepts can effectively leverage shared information to solve related tasks simultaneously. Recent research has shifted focus from specific algorithms or minimization problems to deriving the lowest achievable risk given an MTL data model. The benefits of transfer learning are highlighted through risk bounds, using concepts such as contrasts [23], which study high-dimensional linear regression with auxiliary samples characterized by the sparsity of their contrast vector or transfer distances [29], which provide lower bounds for target generalization

error based on the number of labeled source and target data and their similarities. Additionally, task-correlation matrices [30] have been used to capture the relationships inherent in the data model.

However, these studies remain largely theoretical, lacking practical algorithms to achieve the proposed bounds. They provide broad estimates and orders of magnitude that, while useful for gauging the feasibility of certain objectives, do not offer precise performance evaluations. Moreover, these estimates do not facilitate the optimal tuning of hyperparameters within the MTL framework, a critical aspect of our study. We demonstrate the importance of precise performance evaluations and effective hyperparameter tuning to enhance the practical application of MTL.

Random Matrix Theory in Multi-Task Learning. In the high-dimensional regime, Random Matrix Theory allows obtaining precise theoretical estimates on functionals (e.g., train or test risk) of data matrices with the number of samples comparable to data dimension [2, 13, 31, 46]. In the multi-task classification case, large-dimensional analysis has been performed for Least Squares SVM [48] and supervised PCA [47] showing counterproductivity of tackling tasks independently. We inspire our optimization strategy from [48] and conduct theoretical analysis for the multi-task regression, which introduces unique theoretical challenges and with a different intuition since we need to consider another data generation model and a regression learning objective. Our data modeling shares similarity with [52] which theoretically studied conditions of a negative transfer under covariate and model shifts. Our work focuses on the selection of hyperparameters within a general optimization framework, considering both a hyperparameter λ to relate all tasks and specific parameters γ_t to balance overfitting and underfitting within each task, accommodating non-Gaussian distributions and employing recent developments such as deterministic equivalents and concentration inequalities.

High dimensional analysis for regression. Regression is an important problem that has been extensively explored in the context of single-task learning for high-dimensional data, employing tools such as Random Matrix Theory [12], physical statistics [9, 15], and the Convex Gaussian MinMax Theorem (CGMT) [43], among others. Typically, authors employ a linear signal plus noise model to calculate the asymptotic test risk as a function of the signal-generating parameter and the covariance of the noise. Our research builds upon these studies, extending them to a multi-task learning framework. In doing so, we derive several insights unique to multi-task learning. However, none of these previous studies offer a practical method for estimating the asymptotic test risk, which is dependent on the ground truth signal-generating parameter. Therefore, our work not only extends previous studies into the multi-task case within a generic data distribution (under the assumption of a concentrated random vector), but also provides a practical method for estimating these quantities.

3 Framework

Notation. Throughout this study, matrices are represented by bold uppercase letters (e.g., matrix \mathbf{A}), vectors by bold lowercase letters (e.g., vector \mathbf{v}), and scalars by regular, non-bold typeface (e.g., scalar a). The notation $\mathbf{A} \otimes \mathbf{B}$ for matrices or vectors \mathbf{A}, \mathbf{B} is the Kronecker product. $\mathcal{D}_{\mathbf{x}}$ or $\text{Diag}(\mathbf{x})$ stands for a diagonal matrix containing on its diagonal the elements of the vector \mathbf{x} . The superscripts t and i are used to denote the task and the sample number, respectively, e.g., $\mathbf{x}_i^{(t)}$ writes the i -th sample of the t -th task. The canonical vector of \mathbb{R}^T is denoted by $\mathbf{e}_t^{[T]}$ with $[\mathbf{e}_t^{[T]}]_i = \delta_{ti}$. Given a matrix $M \in \mathbb{R}^{p \times n}$, the Frobenius norm of \mathbf{M} is denoted $\|\mathbf{M}\|_F \equiv \sqrt{\text{tr}(\mathbf{M}^\top \mathbf{M})}$. For our theoretical analysis, we introduce the following notation of training data:

$$\mathbf{Y} = [\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(T)}] \in \mathbb{R}^{q \times n}, \quad \mathbf{Z} = \sum_{t=1}^T \left(\mathbf{e}_t^{[T]} \mathbf{e}_t^{[T]\top} \right) \otimes \mathbf{X}^{(t)} \in \mathbb{R}^{T d \times n}$$

where $n = \sum_{t=1}^T n_t$ is the total number of samples in all the tasks.

3.1 Regression Model

We define $\mathbf{V} = [\mathbf{V}_1^\top, \dots, \mathbf{V}_T^\top]^\top \in \mathbb{R}^{dT \times q}$ and propose to solve the linear multi-task problem by finding $\hat{\mathbf{W}} = [\hat{\mathbf{W}}_1^\top, \dots, \hat{\mathbf{W}}_T^\top]^\top \in \mathbb{R}^{dT \times q}$ which solves the following optimization problem under the assumption of relatedness between the tasks, i.e., $\mathbf{W}_t = \mathbf{W}_0 + \mathbf{V}_t$ for all tasks t :

$$\min_{(\mathbf{W}_0, \mathbf{V}) \in \mathbb{R}^{d \times q} \times \mathbb{R}^{dT \times q}} \mathcal{J}(\mathbf{W}_0, \mathbf{V}) \quad (4)$$

where

$$\mathcal{J}(\mathbf{W}_0, \mathbf{V}) \equiv \frac{1}{2\lambda} \|\mathbf{W}_0\|_F^2 + \frac{1}{2} \sum_{t=1}^T \frac{\|\mathbf{V}_t\|_F^2}{\gamma_t} + \frac{1}{2} \sum_{t=1}^T \left\| \mathbf{Y}^{(t)} - \frac{\mathbf{X}^{(t)\top} \mathbf{W}_t}{\sqrt{Td}} \right\|_F^2.$$

The objective function consists of three components: a regularization term for \mathbf{W}_0 to mitigate overfitting, a task-specific regularization term that controls deviations \mathbf{V}_t from the shared weight matrix \mathbf{W}_0 , and a loss term quantifying the error between the predicted outputs and the actual responses for each task.

The optimization problem is convex for any positive values of $\lambda, \gamma_1, \dots, \gamma_T$, and it possesses a unique solution. The details of the calculus are given in Appendix B. The formula for $\hat{\mathbf{W}}_t$ is given as follows:

$$\hat{\mathbf{W}}_t = \left(\mathbf{e}_t^{[T]\top} \otimes \mathbf{I}_d \right) \frac{\mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Y}}{\sqrt{Td}}, \quad \hat{\mathbf{W}}_0 = (\mathbf{1}_T^\top \otimes \lambda \mathbf{I}_d) \frac{\mathbf{Z} \mathbf{Q} \mathbf{Y}}{\sqrt{Td}},$$

where with $\mathbf{Q} = \left(\frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} + \mathbf{I}_n \right)^{-1} \in \mathbb{R}^{n \times n}$, and $\mathbf{A} = (\mathcal{D}_\gamma + \lambda \mathbf{1}_T \mathbf{1}_T^\top) \otimes \mathbf{I}_d \in \mathbb{R}^{Td \times Td}$.

3.2 Assumptions

In order to use Random Matrix Theory (RMT) tools, we make two assumptions on the data distribution and the asymptotic regime. Following [50], we adopt a concentration hypothesis on the feature vectors $\mathbf{x}_i^{(t)}$, which was shown to be highly effective for analyzing machine learning problems [11, 14]. The justification and implications of these assumptions can be found in appendix A.2 and A.3.

Assumption 1 (Concentrated Random Vector). *We assume that there exists two constants $C, c > 0$ (independent of dimension d) such that, for any task t , for any 1-Lipschitz function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, any feature vector $\mathbf{x}^{(t)} \in \mathcal{X}^{(t)}$ verifies:*

$$\forall t > 0 : \mathbb{P}(|f(\mathbf{x}^{(t)}) - \mathbb{E}[f(\mathbf{x}^{(t)})]| \geq t) \leq C e^{-(t/c)^2},$$

$$\mathbb{E}[\mathbf{x}^{(t)}] = 0 \quad \text{and} \quad \text{Cov}[\mathbf{x}^{(t)}] = \Sigma^{(t)}.$$

In particular, we distinguish the following scenarios: $\mathbf{x}_i^{(t)} \in \mathbb{R}^d$ are concentrated when they are (i) independent Gaussian random vectors with covariance of bounded norm, (ii) independent random vectors uniformly distributed on the \mathbb{R}^d sphere of radius \sqrt{d} , and most importantly (iii) any Lipschitz transformation $\phi(\mathbf{x}_i^{(t)})$ of the above two cases, with bounded Lipschitz norm. Scenario (iii) is especially pertinent for modeling data in realistic settings. Recent research [39] has demonstrated that images produced by generative adversarial networks (GANs) are inherently qualified as concentrated random vectors.

Next, we present a classical RMT assumption that establishes a commensurable relationship between the number of samples and dimension.

Assumption 2 (High-dimensional asymptotics). *As $d \rightarrow \infty$, $n_t = \mathcal{O}(d)$ and $T = \mathcal{O}(1)$. More specifically, we assume that $n/d \xrightarrow{a.s.} c_0 < \infty$ with $n = \sum_{t=1}^T n_t$.*

Although different from classical asymptotic where the number of samples is implicitly assumed to be exponentially larger than the dimension, the high-dimensional asymptotic finds many applications including telecommunications [10], finance [34] and machine learning [11, 14, 48].

4 Main Theoretical Results

4.1 Estimation of the Performances

Given the training dataset $\mathbf{X} \in \mathbb{R}^{n \times d}$ with the corresponding response variable $\mathbf{Y} \in \mathbb{R}^{n \times q}$ and for any test dataset $\mathbf{x}^{(t)} \in \mathbb{R}^d$ and the corresponding response variable $\mathbf{y}^{(t)} \in \mathbb{R}^q$, we aim to derive the theoretical training and test risks defined as follows:

$$\mathcal{R}_{train}^\infty = \frac{1}{Tn} \mathbb{E} [\|\mathbf{Y} - g(\mathbf{X})\|_2^2], \quad \mathcal{R}_{test}^\infty = \frac{1}{T} \sum_{t=1}^T \mathbb{E} [\|\mathbf{y}^{(t)} - g(\mathbf{x}^{(t)})\|_2^2]$$

The output score $g(\mathbf{x}^{(t)}) \in \mathbb{R}^q$ for data $\mathbf{x}^{(t)} \in \mathbb{R}^d$ of task t is given by:

$$g(\mathbf{x}^{(t)}) = \frac{1}{\sqrt{Td}} \left(\mathbf{e}_t^{[T]} \otimes \mathbf{x}^{(t)} \right)^\top \hat{\mathbf{W}}_t = \frac{1}{Td} \left(\mathbf{e}_t^{[T]} \otimes \mathbf{x}^{(t)} \right)^\top \mathbf{AZQY} \quad (5)$$

In particular, the output for the training score is given by:

$$g(\mathbf{X}) = \frac{1}{Td} \mathbf{Z}^\top \mathbf{AZQY}, \quad \mathbf{Q} = \left(\frac{\mathbf{Z}^\top \mathbf{AZ}}{Td} + \mathbf{I}_{Td} \right)^{-1} \quad (6)$$

To understand the statistical properties of $\mathcal{R}_{train}^\infty$ and $\mathcal{R}_{test}^\infty$ for the linear generative model, we study the statistical behavior of the resolvent matrix \mathbf{Q} defined in (6). The notion of a deterministic equivalent from random matrix theory [16] is particularly useful here as it allows us to design a deterministic matrix equivalent to \mathbf{Q} in the sense of linear forms. Specifically, a deterministic equivalent $\bar{\mathbf{M}}$ of a random matrix \mathbf{M} is a deterministic matrix that satisfies $u(\mathbf{M} - \bar{\mathbf{M}}) \rightarrow 0$ almost surely for any bounded linear form $u : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$. We denote this as $\mathbf{M} \leftrightarrow \bar{\mathbf{M}}$. This concept is crucial to our analysis because we need to estimate quantities such as $\frac{1}{d} \text{tr}(\mathbf{AQ})$ or $\mathbf{a}^\top \mathbf{Q} \mathbf{b}$, which are bounded linear forms of \mathbf{Q} with \mathbf{A} , \mathbf{a} and \mathbf{b} all having bounded norms. For convenience, we further express some contributions of the performances with the so-called ‘‘coresolvent’’ defined as $\tilde{\mathbf{Q}} \equiv \left(\frac{\mathbf{A}^{\frac{1}{2}} \mathbf{Z} \mathbf{Z}^\top \mathbf{A}^{\frac{1}{2}}}{Td} + \mathbf{I}_{Td} \right)^{-1}$.

Using Lemma 1 provided in the Appendix C.1, whose proofs are included in Appendices C.2, C.3 and C.4, we establish the deterministic equivalents that allow us to introduce our Theorems 5 and 1, characterizing the asymptotic behavior of both training and testing risks.

Theorem 1 (Asymptotic train and test risk). *Assuming that the training data vectors $\mathbf{x}_i^{(t)}$ and the test data vectors $\mathbf{x}^{(t)}$ are concentrated random vectors, and given the growth rate assumption (Assumption 2), it follows that:*

$$\mathcal{R}_{test}^\infty = \underbrace{\frac{\text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2(\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W} \right)}{Td}}_{\text{signal term}} + \underbrace{\frac{\text{tr}(\Sigma_n \bar{\mathbf{Q}}_2)}{Td} + \text{tr}(\Sigma_n)}_{\text{noise terms}}. \quad (\text{ATR})$$

In addition, the asymptotic risk on the training data is given by

$$\mathcal{R}_{train}^\infty \leftrightarrow \frac{\text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \tilde{\mathbf{Q}} \mathbf{A}^{-1/2} \mathbf{W} \right)}{Tn} - \frac{\text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \tilde{\mathbf{Q}}_2(\mathbf{I}_{Td}) \mathbf{A}^{-1/2} \mathbf{W} \right)}{Tn} + \frac{\text{tr}(\Sigma_n \bar{\mathbf{Q}}_2)}{Tn},$$

where $\tilde{\mathbf{Q}}$, $\tilde{\mathbf{Q}}_2(\mathbf{I}_{Td})$ and $\bar{\mathbf{Q}}_2$ are respectively deterministic equivalents for $\tilde{\mathbf{Q}}$, $\tilde{\mathbf{Q}}^2$ and \mathbf{Q}^2 .

We defer the full proof of this theorem to Appendix D and provide an outline of the test risk proof below. Note that derivations of the asymptotic risk for the training data are more complex and involve computing deterministic equivalents (Lemma 1 of Appendix C.1), which is a powerful tool commonly used in Random Matrix Theory. In the derived theorem, we achieve a stronger convergence result compared to the almost surely convergence. Specifically, we prove a concentration result for the training and test risk with an exponential convergence rate for sufficiently large values of n and d .

$$\begin{aligned}
\mathcal{R}_{test}^\infty &= \frac{1}{T} \mathbb{E}[\|\mathbf{y} - g(\mathbf{x})\|_2^2] \\
&= \frac{1}{T} \mathbb{E} \left[\left\| \frac{\mathbf{x}^\top \mathbf{W}}{\sqrt{Td}} + \varepsilon - \frac{\mathbf{x}^\top \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{W}}{Td\sqrt{Td}} - \frac{\mathbf{x}^\top \mathbf{A} \mathbf{Z} \mathbf{Q} \varepsilon}{Td} \right\|_2^2 \right] \\
&= \frac{1}{T} \mathbb{E} \left[\frac{\text{tr}(\mathbf{W}^\top \boldsymbol{\Sigma} \mathbf{W})}{Td} - \frac{2\text{tr}(\mathbf{W}^\top \boldsymbol{\Sigma} \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{W})}{(Td)^2} + \text{tr}(\varepsilon^\top \varepsilon) + \frac{\text{tr}(\mathbf{W}^\top \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{A} \boldsymbol{\Sigma} \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{W})}{(Td)^3} + \right. \\
&\quad \left. \frac{\text{tr}(\varepsilon^\top \mathbf{Q} \mathbf{Z}^\top \mathbf{A} \boldsymbol{\Sigma} \mathbf{A} \mathbf{Z} \mathbf{Q} \varepsilon)}{(Td)^2} \right] \\
&= \frac{\text{tr}(\mathbf{W}^\top \boldsymbol{\Sigma} \mathbf{W})}{Td} - 2 \frac{\text{tr}(\mathbf{W}^\top \boldsymbol{\Sigma} \mathbf{A}^{\frac{1}{2}} (\mathbf{I}_{Td} - \tilde{\mathbf{Q}}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W})}{Td} + \text{tr}(\varepsilon^\top \varepsilon) + \frac{\text{tr}(\mathbf{W}^\top \boldsymbol{\Sigma} \mathbf{W})}{Td} \\
&\quad - 2 \frac{\text{tr}(\mathbf{W}^\top \boldsymbol{\Sigma} \mathbf{A}^{\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{-\frac{1}{2}} \mathbf{W})}{Td} + \frac{\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2 (\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W}}{Td} + \frac{1}{Td} \text{tr}(\boldsymbol{\Sigma}_N \tilde{\mathbf{Q}}_2) + O\left(\frac{1}{\sqrt{d}}\right)
\end{aligned}$$

4.2 Error Contribution Analysis

To gain theoretical insights, we analyze (ATR) consisting of the signal and the noise components.

Signal Term. The signal term can be further approximated, up to some constants as $\text{tr}(\mathbf{W}^\top (\mathbf{A} \boldsymbol{\Sigma} + \mathbf{I})^{-2} \mathbf{W})$ with $\boldsymbol{\Sigma} = \sum_{t=1}^T \frac{n_t}{d} \boldsymbol{\Sigma}^{(t)}$. The matrix $(\mathbf{A} \boldsymbol{\Sigma} + \mathbf{I})^{-2}$ plays a crucial role in amplifying the signal term $\text{tr}(\mathbf{W}^\top \mathbf{W})$, which in turn allows the test risk to decrease. The off-diagonal elements of $\mathbf{A} \boldsymbol{\Sigma} + \mathbf{I}$ amplify the cross terms ($\text{tr}(\mathbf{W}_v^\top \mathbf{W}_t)$ for $t \neq v$), enhancing the multi-task aspect, while the diagonal elements amplify the independent terms ($\|\mathbf{W}_t\|_2^2$). This structure is significant in determining the effectiveness of multi-task learning.

Furthermore, both terms decrease with an increasing number of samples n_t , smaller values of γ_t , and a larger value of λ . The cross term, which is crucial for multi-task learning, depends on the matrix $\boldsymbol{\Sigma}_t^{-1} \boldsymbol{\Sigma}_v$. This matrix represents the shift in covariates between tasks. When the features are aligned (i.e., $\boldsymbol{\Sigma}_t^{-1} \boldsymbol{\Sigma}_v = \mathbf{I}_d$), the cross term is maximized, enhancing multi-task learning. However, a larger Fisher distance between the covariances of the tasks results in less favorable correlations for multi-task learning.

Noise term. Similar to the signal term, the noise term can be approximated, up to some constants, as $\text{tr}(\boldsymbol{\Sigma}_N (\mathbf{A}^{-1} + \boldsymbol{\Sigma})^{-1})$. However, there is a major difference between the way both terms are expressed in the test risk. The noise term does not include any cross terms because the noise across different tasks is independent. In this context, only the diagonal elements of the matrix are significant. This diagonal term increases with the sample size and the value of λ . It is responsible for what is known as negative transfer. As the diagonal term increases, it negatively affects the transfer of learning from one task to another. This is a critical aspect to consider in multi-task learning scenarios.

4.3 Simplified Model for Clear Insights

In this section, we specialize the theoretical analysis to the simple case of two tasks ($T = 2$). We assume that the tasks share the same identity covariance and that $\gamma_1 = \gamma_2 \equiv \gamma$. Under these conditions, the test risk can be approximated, up to some constants, as

$$\mathcal{R}_{test}^\infty = \mathbf{D}_{IL} (\|\mathbf{W}_1\|_2^2 + \|\mathbf{W}_2\|_2^2) + \mathbf{C}_{MTL} \mathbf{W}_1^\top \mathbf{W}_2 + \mathbf{N}_{NT} \text{tr} \boldsymbol{\Sigma}_n$$

where the diagonal term (independent learning) \mathbf{D}_{IL} , the cross term (multi-task learning) \mathbf{C}_{MTL} , and the noise term (negative transfer) \mathbf{N}_{NT} have closed-form expressions depending on γ and λ :

$$\begin{aligned}
\mathbf{D}_{IL} &= \frac{(c_0(\lambda + \gamma) + 1)^2 + c_0^2 \lambda^2}{(c_0(\lambda + \gamma) + 1)^2 - c_0^2 \lambda^2}, & \mathbf{C}_{MTL} &= \frac{-2c_0 \lambda (c_0(\lambda + \gamma) + 1)}{(c_0(\lambda + \gamma) + 1)^2 - c_0^2 \lambda^2} \\
\mathbf{N}_{NT} &= \frac{(c_0(\lambda + \gamma)^2 + (\lambda + \gamma) - c_0 \lambda^2)^2 + \lambda^2}{((c_0(\lambda + \gamma) + 1)^2 - c_0^2 \lambda^2)^2}
\end{aligned}$$




Figure 1: Test loss contributions D_{IL} , C_{MTL} , N_{NT} across three sample size regimes. Test risk exhibits decreasing, increasing, or convex shapes based on the regime. λ^* from theory are marked.

We recall that c_0 has been defined in the Assumption 2. As previously mentioned, the test risk is primarily composed of two terms: the signal term and the noise term, which are in competition with each other. The more similar the tasks are, the stronger the signal term becomes. In the following plot, we illustrate how this competition can influence the risk. Depending on the value of the parameter λ and the sample sizes, the risk can either decrease monotonically, increase monotonically, or exhibit a convex behavior. This competition can lead to an optimal value for λ , which interestingly has a simple closed-form expression that can be obtained by deriving the $\mathcal{R}_{test}^\infty$ w.r.t. λ as follows (see details in Appendix E.4):

$$\lambda^* = \frac{n}{d} SNR - \frac{\gamma}{2}, \text{ with } SNR = \frac{\|\mathbf{W}_1\|_2^2 + \|\mathbf{W}_2\|_2^2}{\text{tr}\Sigma_N} + \frac{\mathbf{W}_1^\top \mathbf{W}_2}{\text{tr}\Sigma_N}.$$

4.4 Comparison between Empirical and Theoretical Predictions

In this section, we compare the theoretical predictions with the empirical results. Our experiment is based on a two-task setting ($T = 2$) defined as $\mathbf{W}_1 \sim \mathcal{N}(0, I_p)$ with $\mathbf{W}_2 = \alpha \mathbf{W}_1 + \sqrt{1 - \alpha^2} \mathbf{W}_1^\perp$. \mathbf{W}_1^\perp represents any vector orthogonal to \mathbf{W}_1 and $\alpha \in [0, 1]$. This setting allows us to adjust the similarity between tasks through α .

Figure 2 shows a comparison of the theoretical and empirical classification errors for different values of λ , highlighting the error-minimizing value of λ . Despite the relatively small values of n and p , there is a very precise match between the asymptotic theory and the practical experiment. This is particularly evident in the accurate estimation of the optimal value for λ .

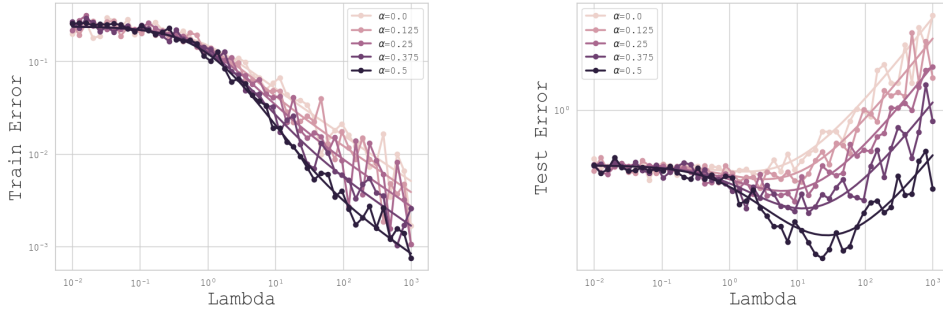


Figure 2: Empirical and theoretical train and test MSE as functions of the parameter λ for different values of α . The smooth curves represent the theoretical predictions, while the corresponding curves with the same color show the empirical results, highlighting that the empirical observations indeed match the theoretical predictions.




Figure 3: Theoretical vs Empirical MSE as function of regularization parameter λ . Close fit between the theoretical and the empirical predictions which underscores the robustness of the theory in light of varying assumptions as well as the accuracy of the suggested estimates. We consider the first two channels as the the two tasks and $d = 144$. 95 samples are used for the training and 42 samples are used for the test.

5 Hyperparameter Optimization in Practice

5.1 Empirical Estimation of Task-wise Signal, Cross Signal, and Noise

All the quantities defined in Theorem 1 are known except for the bilinear expressions $\frac{1}{Td} \text{tr}(\mathbf{W}^\top \mathbf{M} \mathbf{W})$ and $\frac{1}{Td} \text{tr}(\mathbf{\Sigma}_N \mathbf{M})$. These quantities can be consistently estimated under Assumptions 2 as follows (see details in Section F of the Appendix) :

$$\frac{1}{Td} \text{tr}(\mathbf{W}^\top \mathbf{M} \mathbf{W}) - \zeta(\mathbf{M}) \xrightarrow{\text{a.s.}} 0, \quad \frac{1}{Td} \text{tr}(\mathbf{\Sigma}_N \mathbf{M}) - \hat{\sigma} \text{tr} \mathbf{M} \xrightarrow{\text{a.s.}} 0$$

where $\zeta(\mathbf{M}) = \frac{1}{Td} \text{tr}(\hat{\mathbf{W}}^\top \kappa^{-1}(\mathbf{M}) \mathbf{M} \hat{\mathbf{W}}) - \frac{\hat{\sigma}}{Td} \text{tr} \tilde{\mathbf{Q}}_2(\mathbf{A}^{\frac{1}{2}} \kappa^{-1}(\mathbf{M}) \mathbf{A}^{\frac{1}{2}})$.

We define the estimate of the noise as $\hat{\sigma} = \lim_{\substack{\lambda \rightarrow 0 \\ \gamma_t \rightarrow \infty}} \mathcal{R}_{train}^\infty$ and the function κ^{-1} is the functional inverse of the mapping $\kappa : \mathbb{R}^{Td \times Td} \rightarrow \mathbb{R}^{q \times q}$ defined as follows

$$\kappa(\mathbf{M}) = \mathbf{M} - 2\mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{M} + \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2 (\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}}) \mathbf{A}^{-\frac{1}{2}} - \frac{n}{(Td)^2} \mathbf{A}^{-\frac{1}{2}} \mathbf{\Sigma} \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2 (\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}}).$$

5.2 Application to Multi-task Regression

In our study, we apply the theoretical framework presented in our paper to a real-world regression problem, specifically, the *Appliance Energy dataset* which aims to predict the total usage of a house. This dataset is a multivariate regression dataset containing 138 time series each of dimension 24. We select two of these features as 2 tasks to cast the problem as a multi-task learning regression problem.

Figure 3 presents a comparison between the theoretical predictions and empirical outcomes of our proposed linear framework. Despite the assumptions made in the main body of the paper, the theoretical predictions closely align with the experimental results. This demonstrates that our estimates are effective in practice and provide a reliable approximation of the optimal regularization.

In essence, our theoretical framework, when applied to real-world multi-task learning regression problems, yields practical and accurate estimates, thereby validating its effectiveness and applicability.

5.3 Relevance of the theoretical insights beyond the case of linear models

While non-linear models are widely used, establishing their theoretical foundations is challenging. Therefore, we focused on linear models, which, despite their simplicity, provide valuable insights into more complex models.

Our results show that test risk curves for non-linear models follow patterns predicted by our theory. This is expected because non-linear models in time series forecasting typically use a linear output layer for prediction. Thus, we can apply our theory to the inputs of this final linear layer. This approach is valid due to data concentration and the Lipschitz nature of neural networks, ensuring outputs of the non-linear part don't deviate significantly from the inputs.

Moreover, multivariate time series models often treat channels separately using univariate methods, missing cross-channel information. Our results in Section 5.4 show that our method surpasses univariate baselines by optimally regularizing with λ and γ , supporting our theory's applicability to non-linear models as the final linear layer effectively leverages concentrated inputs.

Finally, our regularization approach differs from traditional cross-task regularizations that use one task per dataset. We consider each prediction as a task and introduce γ_t parameters alongside λ . These parameters enforce multivariate regularization and control underfitting or overfitting per task. This method is tractable since it's applied at the model's final layer.

The similarity between curves for non-linear and linear models indicates our findings are robust; non-linear models also exhibit optimal regularization parameters, enhancing performance in multivariate forecasting.

5.4 Application to Multivariate Time Series Forecasting

Our theoretical framework is applied in the context of Multivariate Time Series Forecasting, with related work detailed in Appendix G.1. We previously applied this framework in a linear setting, and now aim to evaluate its empirical validity in the non-linear setting of neural networks. The results presented in this section represent the best test MSE, assuming the ability to find the optimal lambda value, which can be considered as an oracle scenario. A study of these limitations can be found in Appendix H.

Motivation. Our approach is applied to the MTSF setting for several reasons. Firstly, many models currently used are essentially univariate, where predictions for individual series are simply concatenated without exploiting the multivariate information inherent in traditional benchmarks. Given that these benchmarks are designed for multivariate forecasting, leveraging multivariate information should yield better results. Secondly, our theoretical framework can benefit this domain, as most predictive models use a linear layer on top of the model to project historical data of length d for predicting the output of length q . This characteristic aligns well with our method, making it a promising fit for enhancing forecasting accuracy.

Our approach. We propose a novel method for MTSF by modifying the loss function to incorporate both individual feature transformations f_t and a shared transformation f_0 . Each univariate-specific transformation f_t is designed to capture the unique dynamics of its respective feature, while f_0 serves as a common transformation applied across all features to capture underlying patterns shared among them. We consider a neural network f with inputs $\mathbf{X} = [\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \dots, \mathbf{X}^{(T)}]$, where T is the number of channels and $\mathbf{X}^{(t)} \in \mathbb{R}^{n \times d}$. For a univariate model without MTL regularization, we predict $\mathbf{Y} = [\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}, \dots, \mathbf{Y}^{(T)}] = [f_1(\mathbf{X}^{(1)}), \dots, f_T(\mathbf{X}^{(T)})]$ and $\mathbf{Y}^{(t)} \in \mathbb{R}^{n \times q}$. We compare these models with their corresponding versions that include MTL regularization, formulated as: $f_t^{MTL}(\mathbf{X}^{(t)}) = f_t(\mathbf{X}^{(t)}) + f_0(\mathbf{Y})$ with $f_t : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times q}$ and $f_0 : \mathbb{R}^{n \times qT} \rightarrow \mathbb{R}^{n \times q}$. We define our regularized loss as follows:

$$\mathcal{L}(\mathbf{X}, \mathbf{Y}) = \sum_{t=1}^T \|\mathbf{Y}^{(t)} - f_t^{MTL}(\mathbf{X}^{(t)})\|_F^2 + \lambda \|f_0(\mathbf{X})\|_F^2 + \sum_{t=1}^T \gamma_t \|f_t(\mathbf{X}^{(t)})\|_F^2, \quad \forall t \in \{1, \dots, T\}.$$

where $\mathbf{Y}^{(t)}$ are the true predictions, f_t represents the univariate model for each channel t , and λ is our regularization parameter, for which we have established a closed form in the case of linear f_t . f_0 serves a role equivalent to W_0 , which was defined in our theoretical study and allows for the regularization of the common part. This component can be added at the top of a univariate model. The parameters γ_t enable the regularization of the specialized parts f_t .

In our setup, f_t is computed in a similar way as in the model without regularization and f_0 is computed by first flattening the concatenation of the predictions of $\mathbf{X}^{(t)}$, then applying a linear projection leveraging common multivariate information before reshaping. The loss function is specifically designed to balance fitting the multivariate series using f_0 and the specific channels using f_t . This

Table 1: MTL regularization results. Algorithms marked with \dagger are state-of-the-art multivariate models and serve as baseline comparisons. All others are univariate. We compared the models with MTL regularization to their corresponding versions without regularization. Each MSE value is derived from 3 different random seeds. MSE values marked with * indicate that the model with MTL regularization performed significantly better than its version without regularization, according to a Student’s t-test with a p-value of 0.05. MSE values are in **bold** when they are the best in their row, indicating the top-performing models.

Dataset	H	with MTL regularization			without MTL regularization					
		PatchTST	DLinearU	Transformer	PatchTST	DLinearU	DLinearM	Transformer	SAMformer \dagger	iTransformer \dagger
ETTh1	96	0.385	0.367*	0.368	0.387	0.397	0.386	0.370	0.381	0.386
	192	0.422	0.405*	0.407*	0.424	0.422	0.437	0.411	0.409	0.441
	336	0.433*	0.431	0.433	0.442	0.431	0.481	0.437	0.423	0.487
	720	0.430*	0.454	0.455*	0.451	0.428	0.519	0.470	0.427	0.503
ETTh2	96	0.291	0.267*	0.270	0.295	0.294	0.333	0.273	0.295	0.297
	192	0.346*	0.331*	0.337	0.351	0.361	0.477	0.339	0.340	0.380
	336	0.332*	0.367	0.366*	0.342	0.361	0.594	0.369	0.350	0.428
	720	0.384*	0.412	0.405*	0.393	0.395	0.831	0.428	0.391	0.427
Weather	96	0.148	0.149*	0.154*	0.149	0.196	0.196	0.170	0.197	0.174
	192	0.190	0.206*	0.198*	0.193	0.243	0.237	0.214	0.235	0.221
	336	0.242*	0.249*	0.258	0.246	0.283	0.283	0.260	0.276	0.278
	720	0.316*	0.326*	0.331	0.322	0.339	0.345	0.326	0.334	0.358

approach enhances the model’s generalization across various forecasting horizons and datasets. More details on our regularized loss function can be found in Appendix G.2

Results. We present experimental results on different forecasting horizons, using common benchmark MTSF datasets, the characteristics of which are outlined in the Appendix G.3. Our models include PatchTST [32], known to be on par with state-of-the-art in MTSF while being a univariate model, a univariate DLinear version called DLinearU compared to its multivariate counterpart DLinearM [53], and a univariate Transformer [19] with temporal-wise attention compared to the multivariate state-of-the-art models SAMformer [19] and iTransformer [24]. Table 1 provides a detailed comparison of the test mean squared errors (MSE) for different MTSF models, emphasizing the impact of MTL regularization. Models with MTL regularization are compared to their versions without regularization, as well as SAMformer and iTransformer. All the experiments can be found in Appendix G.4.

Adding MTL regularization improves the performance of PatchTST, DLinearU, and Transformer in most cases. When compared to state-of-the-art multivariate models, the MTL-regularized models are often competitive. SAMformer is outperformed by at least one MTL-regularized method per horizon and dataset, except for ETTh1 with horizons of 336 and 720. iTransformer is consistently outperformed by at least one MTL-regularized methods regardless of the dataset and horizon.

The best performing methods are PatchTST and DLinearU with MTL regularization. These models not only outperform their non-regularized counterparts, often significantly as shown by Student’s t-tests with a p-value of 0.05, but also surpass state-of-the-art multivariate models like SAMformer and Transformer. This superior performance is indicated by the bold values in the table.

Finally, MTL regularization enhances the performance of univariate models, making them often competitive with state-of-the-art multivariate methods like SAMformer and iTransformer. This approach seems to better captures shared dynamics among tasks, leading to more accurate forecasts.

6 Conclusions and Future Works

In this article, we have explored linear multi-task learning by deriving a closed-form solution for an optimization problem that capitalizes on the wealth of information available across multiple tasks. Leveraging Random Matrix Theory, we have been able to obtain the asymptotic training and testing risks, and have proposed several insights into high-dimensional multi-task learning regression. Our theoretical analysis, though based on a simplified model, has been effectively applied to multi-task regression and multivariate forecasting, using both synthetic and real-world datasets. We believe that our work lays a solid foundation for future research, paving the way for using random matrix theory with more complex models, such as deep neural networks, within the multi-task learning framework.

Acknowledgements

Work partially funded by EU project AI4Europe (101070000).

References

- [1] Ando, R. K., Zhang, T., and Bartlett, P. (2005). A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of machine learning research*, 6(11).
- [2] Bai, Z. and Silverstein, J. W. (2010). *Spectral Analysis of Large Dimensional Random Matrices*. Springer Series in Statistics. Springer New York, NY.
- [3] Box, G. E. P. and Jenkins, G. (1990). *Time Series Analysis, Forecasting and Control*. Holden-Day, Inc., USA.
- [4] Box, G. E. P., Jenkins, G. M., and MacGregor, J. F. (1974). Some Recent Advances in Forecasting and Control. *Journal of the Royal Statistical Society Series C*, 23(2):158–179.
- [5] Caruana, R. (1997). Multitask learning. *Machine learning*, 28:41–75.
- [6] Čepulionis, P. and Lukoševičiūtė, K. (2016). Electrocardiogram time series forecasting and optimization using ant colony optimization algorithm. *Mathematical Models in Engineering*, 2(1):69–77.
- [7] Chen, R. and Tao, M. (2021). Data-driven prediction of general hamiltonian dynamics via learning exactly-symplectic maps. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 1717–1727. PMLR.
- [8] Chen, S.-A., Li, C.-L., Arik, S. O., Yoder, N. C., and Pfister, T. (2023). TSMixer: An all-MLP architecture for time series forecasting. *Transactions on Machine Learning Research*.
- [9] Clarté, L., Loureiro, B., Krzakala, F., and Zdeborová, L. (2023). Theoretical characterization of uncertainty in high-dimensional linear classification. *Machine Learning: Science and Technology*, 4(2):025029.
- [10] Couillet, R. and Debbah, M. (2011). *Random matrix methods for wireless communications*. Cambridge University Press.
- [11] Couillet, R. and Liao, Z. (2022). *Random matrix methods for machine learning*. Cambridge University Press.
- [12] Dobriban, E. and Wager, S. (2018). High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279.
- [13] Erdős, L. and Yau, H.-T. (2017). *A Dynamical Approach to Random Matrix Theory*, volume 28 of *Courant Lecture Notes*. American Mathematical Society.
- [14] Feofanov, V., Tiomoko, M., and Virmaux, A. (2023). Random matrix analysis to balance between supervised and unsupervised learning under the low density separation assumption. In *Proceedings of the 40th International Conference on Machine Learning*, pages 10008–10033.
- [15] Gerbelot, C., Abbara, A., and Krzakala, F. (2022). Asymptotic errors for teacher-student convex generalized linear models (or: How to prove kabashima’s replica formula). *IEEE Transactions on Information Theory*, 69(3):1824–1852.
- [16] Hachem, W., Loubaton, P., and Najim, J. (2007). Deterministic equivalents for certain functionals of large random matrices.
- [17] Hu, Y., Li, M., Lu, Q., Weng, H., Wang, J., Zekavat, S. M., Yu, Z., Li, B., Gu, J., Muchnik, S., et al. (2019). A statistical framework for cross-tissue transcriptome-wide association analysis. *Nature genetics*, 51(3):568–576.

- [18] Ilbert, R., Hoang, T. V., and Zhang, Z. (2024a). Data augmentation for multivariate time series classification: An experimental study. In *2024 IEEE 40th International Conference on Data Engineering Workshops (ICDEW)*, pages 128–139.
- [19] Ilbert, R., Odonnat, A., Feofanov, V., Virmaux, A., Paolo, G., Palpanas, T., and Redko, I. (2024b). Unlocking the potential of transformers in time series forecasting with sharpness-aware minimization and channel-wise attention.
- [20] Ilbert, R., Tiomoko, M., Louart, C., Feofanov, V., Palpanas, T., and Redko, I. (2024c). Enhancing multivariate time series forecasting via multi-task learning and random matrix theory. In *NeurIPS Workshop on Time Series in the Age of Large Models*.
- [21] Kim, T., Kim, J., Tae, Y., Park, C., Choi, J.-H., and Choo, J. (2021). Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International Conference on Learning Representations*.
- [22] Kingma, D. and Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, San Diego, CA, USA.
- [23] Li, S., Cai, T. T., and Li, H. (2022). Transfer learning for high-dimensional linear regression: Prediction, estimation and minimax optimality. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):149–173.
- [24] Liu, Y., Hu, T., Zhang, H., Wu, H., Wang, S., Ma, L., and Long, M. (2024). itransformer: Inverted transformers are effective for time series forecasting. In *The Twelfth International Conference on Learning Representations*.
- [25] Louart, C. and Couillet, R. (2021). Spectral properties of sample covariance matrices arising from random matrices with independent non identically distributed columns. *arXiv preprint arXiv:2109.02644*.
- [26] Lounici, K., Pontil, M., Tsybakov, A. B., and Van De Geer, S. (2009). Taking advantage of sparsity in multi-task learning. *arXiv preprint arXiv:0903.1468*.
- [27] Max Planck Institute (n.d.). Weather dataset.
- [28] Mei, S., Fei, W., and Zhou, S. (2011). Gene ontology based transfer learning for protein subcellular localization. *BMC bioinformatics*, 12:1–12.
- [29] Mousavi Kalan, M., Fabian, Z., Avestimehr, S., and Soltanolkotabi, M. (2020). Minimax lower bounds for transfer learning with linear and one-hidden layer neural networks. *Advances in Neural Information Processing Systems*, 33:1959–1969.
- [30] Nguyen, M.-T. and Couillet, R. (2023). Asymptotic bayes risk of semi-supervised multitask learning on gaussian mixture. In *International Conference on Artificial Intelligence and Statistics*, pages 5063–5078. PMLR.
- [31] Nica, A. and Speicher, R. (2006). *Lectures on the Combinatorics of Free Probability*. London Mathematical Society Lecture Note Series. Cambridge University Press.
- [32] Nie, Y., Nguyen, N. H., Sinthong, P., and Kalagnanam, J. (2023). A time series is worth 64 words: Long-term forecasting with transformers. In *The Eleventh International Conference on Learning Representations*.
- [33] Niu, S., Liu, Y., Wang, J., and Song, H. (2020). A decade survey of transfer learning (2010–2020). *IEEE Transactions on Artificial Intelligence*, 1(2):151–166.
- [34] Potters, M., Bouchaud, J.-P., and Laloux, L. (2005). Financial applications of random matrix theory: Old laces and new pieces. *Acta Physica Polonica B*, 36(9):2767.
- [35] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.

- [36] Romera-Paredes, B., Aung, H., Bianchi-Berthouze, N., and Pontil, M. (2013). Multilinear multitask learning. In *International Conference on Machine Learning*, pages 1444–1452. PMLR.
- [37] Ruder, S., Peters, M. E., Swayamdipta, S., and Wolf, T. (2019). Evolution of transfer learning in natural language processing. *arXiv preprint arXiv:1910.07370*.
- [38] Salinas, D., Flunkert, V., Gasthaus, J., and Januschowski, T. (2020). Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191.
- [39] Seddik, M. E. A., Louart, C., Tamaazousti, M., and Couillet, R. (2020). Random matrix theory proves that deep learning representations of gan-data behave as gaussian mixtures. In *International Conference on Machine Learning*.
- [40] Sen, R., Yu, H.-F., and Dhillon, I. (2019). Think globally, act locally: a deep neural network approach to high-dimensional time series forecasting. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA. Curran Associates Inc.
- [41] Shao, L. (2015). Transfer learning for visual categorization: a survey. *IEEE Transactions on Neural Networks and Learning Systems*, 26(5):1019–1034.
- [42] Shin, H.-C., Roth, H. R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., and Summers, R. M. (2016). Deep convolutional neural networks for computer-aided detection: Cnn architectures, dataset characteristics and transfer learning. *IEEE transactions on medical imaging*, 35(5):1285–1298.
- [43] Sifaou, H., Kammoun, A., and Alouini, M.-S. (2021). A precise performance analysis of support vector regression. In *International Conference on Machine Learning*, pages 9671–9680. PMLR.
- [44] Sonkavde, G., Dharrao, D. S., Bongale, A. M., Deokate, S. T., Doreswamy, D., and Bhat, S. K. (2023). Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications. *International Journal of Financial Studies*, 11(3).
- [45] Sorjamaa, A., Hao, J., Reyhani, N., Ji, Y., and Lendasse, A. (2007). Methodology for long-term prediction of time series. *Neurocomputing*, 70(16):2861–2869. Neural Network Applications in Electrical Engineering Selected papers from the 3rd International Work-Conference on Artificial Neural Networks (IWANN 2005).
- [46] Tao, T. (2012). *Topics in Random Matrix Theory*, volume 132 of *Graduate Studies in Mathematics*. American Mathematical Society.
- [47] Tiomoko, M., Couillet, R., and Pascal, F. (2023). PCA-based multi-task learning: a random matrix approach. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 34280–34300. PMLR.
- [48] Tiomoko, M., Couillet, R., and Tiomoko, H. (2020). Large dimensional analysis and improvement of multi task learning.
- [49] UCI (n.d.). Electricity dataset.
- [50] Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press.
- [51] Wu, H., Xu, J., Wang, J., and Long, M. (2021). Autoformer: Decomposition transformers with Auto-Correlation for long-term series forecasting. In *Advances in Neural Information Processing Systems*.
- [52] Yang, F., Zhang, H. R., Wu, S., Ré, C., and Su, W. J. (2023). Precise high-dimensional asymptotics for quantifying heterogeneous transfers.
- [53] Zeng, A., Chen, M., Zhang, L., and Xu, Q. (2023). Are transformers effective for time series forecasting? In *Proceedings of the AAAI Conference on Artificial Intelligence*.

- [54] Zhang, H., Zheng, Y., and Qi, D. (2019). Multi-task learning for time series forecasting using neural networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1305–1313. ACM.
- [55] Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., and Zhang, W. (2021). Informer: Beyond efficient transformer for long sequence time-series forecasting. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual Conference*, volume 35, pages 11106–11115. AAAI Press.
- [56] Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L., and Jin, R. (2022). FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *Proc. 39th International Conference on Machine Learning (ICML 2022)*.

Appendix

Roadmap. This appendix provides the technical details omitted in the main paper. It starts with an overview of the setup considered in the paper in **Section A**. **Section B** offers a detail computation for $\hat{\mathbf{W}}_t$ and $\hat{\mathbf{W}}_0$. **Section C** contains a proof of Lemma 1. **Section D** explains the theoretical steps for deriving the training and test risks, as well as the deterministic equivalents. **Section E** discusses the technical tools used to derive the main intuitions presented by the theory. **Section F** focuses on the derivation of the estimations of the main quantities involved in the training and test risks. **Section G** complements the experimental study by presenting additional experiments. Finally, **Section H** deals with the limitations of our approach in a non-linear setting.

Table of Contents

A Multi-task Learning setup	16
A.1 On the zero-mean assumption	16
A.2 On the Assumption 1	16
A.3 On the Assumption 2	16
B Minimization Problem	16
B.1 Computation of $\hat{\mathbf{W}}_t$ and $\hat{\mathbf{W}}_0$	16
C Lemma 1 and proof with Random Matrix Theory	18
C.1 Lemma 1	18
C.2 Deterministic equivalent of the resolvent $\tilde{\mathbf{Q}}$	18
C.3 Deterministic equivalent of bilinear forms of the resolvent	19
C.4 Estimation of the deterministic equivalent of \mathbf{Q}^2	21
D Risk Estimation (Proof of Theorem 1)	22
D.1 Test Risk	22
D.2 Train Risk	23
E Interpretation and insights of the theoretical analysis	23
E.1 Analysis of the test risk	23
E.2 Interpretation of the signal term	24
E.3 Interpretation and insights of the noise terms	24
E.4 Optimal Lambda	24
F Theoretical Estimations	25
F.1 Estimation of the training and test risk	25
F.2 Estimation of the noise covariance	25
G Multivariate Time Series Forecasting	25
G.1 Related Work	25
G.2 Architecture and Training Parameters	26
G.3 Datasets	26

A Multi-task Learning setup

A.1 On the zero-mean assumption

We would like to note that we are assuming that both the noise ε and the feature $\mathbf{x}_i^{(t)}$ have zero mean. This is a common assumption in many statistical models and it simplifies the analysis. However, this assumption is not restrictive. In practice, if the data or the response variable are not centered, we can always preprocess the data by subtracting the mean. This preprocessing step brings us back to the zero-mean setting that we consider in our theoretical analysis.

A.2 On the Assumption 1

Data are concentrated random vectors, meaning high-dimensional data maintain stable under complex (Lipschitz) transformations. The strong performance of neural networks on tasks like image recognition and NLP suggests that these models produce stable predictions. As Lipschitz transformations, they maintain controlled distances between inputs, ensuring stability. Recent studies have demonstrated and experimentally confirmed that both real-world data and synthetically generated data using GANs exhibit concentration properties, supporting this assumption. This makes our assumption more realistic than traditional Gaussian assumptions, as it does not rely on specific hypotheses about the shape of the data distribution, but rather on the stability of statistical properties after transformation. Consequently, analyzing a framework of concentrated random vectors is more theoretically challenging than using Gaussian assumptions and represents a key novelty of our theory.

A.3 On the Assumption 2

The dimension d is of the same order of magnitude as the sample size n . This joint growth captures data complexity better than assuming a fixed feature size with increasing samples, which can oversimplify models. Our theory works for fixed d and n , unbiased by specific parameter choices. The accuracy of empirical predictions depends on both d and n , with variance scaling as $O\left(\frac{1}{\sqrt{dn}}\right)$. Larger d and n reduce variance, making empirical results more reliable and closer to theoretical values. Conversely, smaller d and n increase variance, affecting single predictions. However, this scaling is still better than n growing indefinitely with fixed d , where variance scales as $O\left(\frac{1}{\sqrt{n}}\right)$, leading to increased bias. Implication of the assumptions. In Section 5.2, Figure 3, the Appliance Energy dataset illustrates our assumptions' realism. Despite the moderate sample size of 42 samples with 142 dimensions and non-synthetic data, the theoretical curve fits the empirical predictions well. Additionally, with synthetic data ($d = 100, n = 100$), the theory matches the empirical curve across various hyperparameters (Figure 2), showcasing the predictive power of Random Matrix Theory.

B Minimization Problem

B.1 Computation of $\hat{\mathbf{W}}_t$ and $\hat{\mathbf{W}}_0$

The proposed multi task regression finds $\hat{\mathbf{W}} = [\hat{\mathbf{W}}_1^\top, \dots, \hat{\mathbf{W}}_k^\top]^\top \in \mathbb{R}^{dT \times q}$ which solves the following optimization problem using the additional assumption of relatedness between the tasks ($\mathbf{W}_t = \mathbf{W}_0 + \mathbf{V}_t$ for all tasks t):

$$\min_{(\mathbf{W}_0, \mathbf{V}) \in \mathbb{R}^d \times \mathbb{R}^{d \times T} \times \mathbb{R}^T} \mathcal{J}(\mathbf{W}_0, \mathbf{V}) \tag{7}$$

where

$$\mathcal{J}(\mathbf{W}_0, \mathbf{V}) \equiv \frac{1}{2\lambda} \text{tr}(\mathbf{W}_0^\top \mathbf{W}_0) + \frac{1}{2} \sum_{t=1}^T \frac{\text{tr}(\mathbf{V}_t^\top \mathbf{V}_t)}{\gamma_t} + \frac{1}{2} \sum_{t=1}^T \text{tr}(\boldsymbol{\xi}_t^\top \boldsymbol{\xi}_t)$$

$$\boldsymbol{\xi}_t = \mathbf{Y}^{(t)} - \frac{\mathbf{X}^{(t)\top} \mathbf{W}_t}{\sqrt{Td}}, \quad \forall t \in \{1, \dots, T\}.$$

The Lagrangian introducing the lagrangian parameters for each task t , $\boldsymbol{\alpha}_t \in \mathbb{R}^{n_t \times q}$ reads as

$$\mathcal{L}(\mathbf{W}_0, \mathbf{V}_t, \boldsymbol{\xi}_t, \boldsymbol{\alpha}_t) = \frac{1}{2\lambda} \text{tr}(\mathbf{W}_0^\top \mathbf{W}_0) + \frac{1}{2} \sum_{t=1}^T \frac{\text{tr}(\mathbf{V}_t^\top \mathbf{V}_t)}{\gamma_t} + \frac{1}{2} \sum_{t=1}^T \text{tr}(\boldsymbol{\xi}_t^\top \boldsymbol{\xi}_t)$$

$$+ \sum_{t=1}^T \text{tr} \left(\boldsymbol{\alpha}_t^\top \left(\mathbf{Y}^{(t)} - \frac{\mathbf{X}^{(t)\top} (\mathbf{W}_0 + \mathbf{V}_t)}{\sqrt{Td}} - \boldsymbol{\xi}_t \right) \right)$$

Differentiating with respect to the unknown variables $\hat{\mathbf{W}}_0$, $\hat{\mathbf{V}}_t$, $\boldsymbol{\xi}_t$, $\boldsymbol{\alpha}_t$ and \mathbf{b}_t , we get the following system of equation

$$\frac{1}{\lambda} \hat{\mathbf{W}}_0 - \sum_{t=1}^T \frac{\mathbf{X}^{(t)} \boldsymbol{\alpha}_t}{\sqrt{Td}} = 0$$

$$\frac{1}{\gamma_t} \hat{\mathbf{V}}_t - \frac{\mathbf{X}^{(t)} \boldsymbol{\alpha}_t}{\sqrt{Td}} = 0$$

$$\boldsymbol{\xi}_t - \boldsymbol{\alpha}_t = 0$$

$$\mathbf{Y}^{(t)} - \frac{\mathbf{X}^{(t)\top} \hat{\mathbf{W}}_0}{\sqrt{Td}} - \frac{\mathbf{X}^{(t)\top} \hat{\mathbf{V}}_t}{\sqrt{Td}} - \boldsymbol{\xi}_t = 0$$

Plugging the expression of $\hat{\mathbf{W}}_0$, $\hat{\mathbf{V}}_t$ and $\boldsymbol{\xi}_t$ into the expression of $\mathbf{Y}^{(t)}$ gives

$$\mathbf{Y}^{(t)} = \lambda \sum_{t=1}^T \frac{\mathbf{X}^{(t)\top} \mathbf{X}^{(t)}}{Td} \boldsymbol{\alpha}_t + \gamma_t \frac{\mathbf{X}^{(t)\top} \mathbf{X}^{(t)}}{Td} \boldsymbol{\alpha}_t + \boldsymbol{\alpha}_t$$

which can be rewritten as

$$\mathbf{Y}^{(t)} = (\lambda + \gamma_t) \frac{\mathbf{X}^{(t)\top} \mathbf{X}^{(t)}}{Td} \boldsymbol{\alpha}_t + \lambda \sum_{v \neq t} \frac{\mathbf{X}^{(t)\top} \mathbf{X}^{(v)}}{Td} \boldsymbol{\alpha}_v + \boldsymbol{\alpha}_t$$

With $\mathbf{Y} = [\mathbf{Y}^{(1)\top}, \dots, \mathbf{Y}^{(T)\top}]^\top \in \mathbb{R}^{n \times q}$, $\boldsymbol{\alpha} = [\boldsymbol{\alpha}_1^\top, \dots, \boldsymbol{\alpha}_k^\top]^\top \in \mathbb{R}^{n \times q}$, $\mathbf{Z} = \sum_{t=1}^T \mathbf{e}_t^{[T]} \mathbf{e}_t^{[T]\top} \otimes \mathbf{X}^{(t)} \in \mathbb{R}^{Td \times n}$, this system of equations can be written under the following compact matrix form:

$$\mathbf{Q}^{-1} \boldsymbol{\alpha} = \mathbf{Y}$$

with $\mathbf{Q} = \left(\frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} + \mathbf{I}_n \right)^{-1} \in \mathbb{R}^{n \times n}$, and $\mathbf{A} = (\mathcal{D}_\gamma + \lambda \mathbf{1}_T \mathbf{1}_T^\top) \otimes \mathbf{I}_d \in \mathbb{R}^{Td \times Td}$.

Solving for $\boldsymbol{\alpha}$ then gives:

$$\boldsymbol{\alpha} = \mathbf{Q} \mathbf{Y}$$

Moreover, using $\hat{\mathbf{W}}_t = \hat{\mathbf{W}}_0 + \hat{\mathbf{V}}_t$, the expression of \mathbf{W}_t becomes:

$$\hat{\mathbf{W}}_t = \left(\mathbf{e}_t^{[T]\top} \otimes \mathbf{I}_d \right) \frac{\mathbf{A} \mathbf{Z} \boldsymbol{\alpha}}{\sqrt{Td}},$$

$$\hat{\mathbf{W}}_0 = (\mathbf{1}_T^\top \otimes \lambda \mathbf{I}_d) \frac{\mathbf{Z} \boldsymbol{\alpha}}{\sqrt{Td}}.$$

C Lemma 1 and proof with Random Matrix Theory

C.1 Lemma 1

Lemma 1 (Deterministic equivalents for $\tilde{\mathbf{Q}}$, $\tilde{\mathbf{Q}}\mathbf{M}\tilde{\mathbf{Q}}$ and \mathbf{Q}^2 for any $\mathbf{M} \in \mathbb{R}^{n \times n}$). *Under the concentrated random vector assumption for each feature vector $\mathbf{x}_i^{(t)}$ and under the growth rate assumption (Assumption 2), for any deterministic $\mathbf{M} \in \mathbb{R}^{n \times n}$, we have the following convergence:*

$$\tilde{\mathbf{Q}} \leftrightarrow \bar{\mathbf{Q}}, \quad \tilde{\mathbf{Q}}\mathbf{M}\tilde{\mathbf{Q}} \leftrightarrow \bar{\mathbf{Q}}_2(\mathbf{M}), \quad \mathbf{Q}^2 \leftrightarrow \bar{\mathbf{Q}}_2$$

where $\bar{\mathbf{Q}}_2$, $\bar{\mathbf{Q}}$ and $\bar{\mathbf{Q}}_2$ are defined as follows

$$\bar{\mathbf{Q}} = \left(\sum_{t=1}^T \frac{c_0 \mathbf{C}^{(t)}}{1 + \delta_t} + \mathbf{I}_{Td} \right)^{-1}, \quad \delta_t = \frac{1}{Td} \text{tr} \left(\boldsymbol{\Sigma}^{(t)} \bar{\mathbf{Q}} \right), \quad \mathbf{C}^{(t)} = \mathbf{A}^{\frac{1}{2}} \left(\mathbf{e}_t^{[T]} \otimes \boldsymbol{\Sigma}^{(t)} \right) \mathbf{A}^{\frac{1}{2}}$$

$$\bar{\mathbf{Q}}_2(\mathbf{M}) = \bar{\mathbf{Q}}\mathbf{M}\bar{\mathbf{Q}} + \frac{1}{Td} \sum_{t=1}^T \frac{d_t}{1 + \delta_t} \bar{\mathbf{Q}}\mathbf{C}^{(t)}\bar{\mathbf{Q}}, \quad \mathbf{d} = \left(\mathbf{I}_T - \frac{1}{Td} \Psi \right)^{-1} \Psi(\mathbf{M}) \in \mathbb{R}^T$$

$$\bar{\mathbf{Q}}_2 = \mathbf{I}_n - \text{Diag}_{t \in [T]}(v_t \mathbf{I}_{n_t}), \quad v_t = \frac{1}{Td} \frac{\text{tr}(\mathbf{C}^{(t)} \bar{\mathbf{Q}})}{(1 + \delta_t)^2} + \frac{1}{Td} \frac{\text{tr}(\mathbf{C}^{(t)} \bar{\mathbf{Q}}_2(\mathbf{I}_n))}{(1 + \delta_t)^2}$$

where

$$\Psi(M) = \left(\frac{n_t}{Td} \frac{\text{tr}(\mathbf{C}^{(t)} \bar{\mathbf{Q}}\mathbf{M}\bar{\mathbf{Q}})}{1 + \delta_t} \right)_{t \in [T]} \in \mathbb{R}^T, \quad \Psi = \left(\frac{n_t}{Td} \frac{\text{tr}(\mathbf{C}^{(t)} \bar{\mathbf{Q}}\mathbf{C}^{(t')} \bar{\mathbf{Q}})}{(1 + \delta_t)(1 + \delta_{t'})} \right)_{t, t' \in [T]} \in \mathbb{R}^{T \times T},$$

C.2 Deterministic equivalent of the resolvent $\tilde{\mathbf{Q}}$

The evaluation of the expectation of linear forms on $\tilde{\mathbf{Q}}$ and $\tilde{\mathbf{Q}}^2$ can be found in the literature. To find a result that meets exactly our setting, we will cite [25] that is a bit more general since it treats cases where $\mathbb{E}[x_i^{(t)}] \neq 0$ for $t \in [T]$ and $i \in [n_t]$. Unlike the main paper, and to be more general, the study presented below is ‘‘quasi asymptotic’’ meaning that the results are true for finite value of d, n . Let us first rewrite the general required hypotheses, adapting them to our setting. For that purpose, we consider in the rest of this paper a certain asymptotic $I \subset \{(d, n), d \in \mathbb{N}, n \in \mathbb{N}\} = \mathbb{N}^2$ satisfying:

$$\{d, \exists n \in \mathbb{N} : (d, n) \in I\} = \mathbb{N} \quad \text{and} \quad \{n, \exists d \in \mathbb{N} : (d, n) \in I\} = \mathbb{N}.$$

such that n and d can tend to ∞ but with some constraint that is given in the first item of Assumption 3 below. Given two sequences $(a_{d,n})_{d,n \in I}, (b_{d,n})_{d,n \in I} > 0$, the notation $a_{d,n} \leq O(b_{d,n})$ (or $a \leq O(b)$) means that there exists a constant $C > 0$ such that for all $(d, n) \in I$, $a_{d,n} \leq C b_{d,n}$.

Assumption 3. *There exists some constants $C, c > 0$ independent such that:*

- $n \leq O(d)$
- $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_n) \in \mathbb{R}^{Td \times n}$ has independent columns
- for any $(d, n) \in I$, and any $f : \mathbb{R}^{Td \times n} \rightarrow \mathbb{R}$ 1-Lipschitz for the euclidean norm:

$$\mathbb{P}(|f(\mathbf{Z}) - \mathbb{E}[f(\mathbf{Z})]| \geq t) \leq C e^{-ct^2}.$$

- $\forall i \in \{n, \exists d \in \mathbb{N}, (d, n) \in I\} : \|\mathbb{E}[\mathbf{z}_i]\| \leq O(1)$.

Theorem 2 ([25], Theorem 0.9.). *Given $T \in \mathbb{N}$, $\mathbf{Z} \in \mathbb{R}^{Td \times n}$ and two deterministic $A \in \mathbb{R}^{Td \times Td}$, we note $\tilde{\mathbf{Q}} \equiv \left(\frac{1}{Td} \mathbf{A}^{\frac{1}{2}} \mathbf{Z} \mathbf{Z}^T \mathbf{A}^{\frac{1}{2}} + \mathbf{I}_{Td} \right)^{-1}$. If \mathbf{Z} satisfies Assumption 3 and $\mathbf{M} \in \mathbb{R}^{Td \times Td}$ is a deterministic matrix satisfying $\|\mathbf{M}\|_F \leq 1$, one has the concentration:*

$$\mathbb{P} \left(\left| \text{tr}(\mathbf{M} \tilde{\mathbf{Q}}) - \text{tr}(\mathbf{M} \bar{\mathbf{Q}}_{\delta(\mathbf{S})}(\mathbf{S})) \right| \geq t \right) \leq C e^{-ct^2},$$

where $\mathbf{S} = (\mathbf{S}_1, \dots, \mathbf{S}_n) = (\mathbb{E}[\mathbf{z}_1 \mathbf{z}_1^\top], \dots, \mathbb{E}[\mathbf{z}_n \mathbf{z}_n^\top])$, for $\boldsymbol{\delta} \in \mathbb{R}^n$, $\bar{\bar{\mathbf{Q}}}_\delta$ is defined as:

$$\bar{\bar{\mathbf{Q}}}_\delta(\mathbf{S}) = \left(\frac{1}{Td} \sum_{i \in [n]} \frac{\mathbf{A}^{\frac{1}{2}} \mathbf{S}_i \mathbf{A}^{\frac{1}{2}}}{1 + \delta_i} + I_{Td} \right)^{-1},$$

and $\boldsymbol{\delta}(\mathbf{S})$ is the unique solution to the system of equations:

$$\forall i \in [n]: \quad \boldsymbol{\delta}(\mathbf{S})_i = \frac{1}{n} \text{tr} \left(\mathbf{A}^{\frac{1}{2}} \mathbf{S}_i \mathbf{A}^{\frac{1}{2}} \bar{\bar{\mathbf{Q}}}_{\boldsymbol{\delta}(\mathbf{S})} \right).$$

We end this subsection with some results that will be useful for next subsection on the estimation of bilinear forms on $\bar{\bar{\mathbf{Q}}}$.

Lemma 2 ([25], Lemmas 4.2, 4.6). *Under the setting of Theorem 2, given a deterministic vector $\mathbf{u} \in \mathbb{R}^{Td}$ such that $\|\mathbf{u}\| \leq O(1)$ and two deterministic matrices \mathbf{U}, \mathbf{V} such that $\|\mathbf{U}\|, \|\mathbf{V}\| \leq O(1)$ and a power $r > 0$, $r \leq O(1)$:*

- $\mathbb{E} \left[\left| \mathbf{u}^\top \mathbf{U} \bar{\bar{\mathbf{Q}}}_{-i} \mathbf{V} \mathbf{z}_i \right|^r \right] \leq O(1)$
- $\mathbb{E} \left[\left| \frac{1}{Td} \mathbf{z}_i^\top \mathbf{U} \bar{\bar{\mathbf{Q}}}_{-i} \mathbf{V} \mathbf{z}_i - \mathbb{E} \left[\frac{1}{Td} \text{tr} \left(\Sigma_i \mathbf{U} \bar{\bar{\mathbf{Q}}}_\mathbf{B} \right) \right] \right|^r \right] \leq O \left(\frac{1}{d^{\frac{r}{2}}} \right).$

C.3 Deterministic equivalent of bilinear forms of the resolvent

To simplify the expression of the following theorem, we take $\mathbf{A} = I_{Td}$. One can replace \mathbf{Z} with $\mathbf{A}^{\frac{1}{2}} \mathbf{Z}$ to retrieve the result necessary for the main paper.

Theorem 3. *Under the setting of Theorem 2, with $\mathbf{A} = I_{Td}$, one can estimate for any deterministic matrices $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{Td}$ such that $\|\mathbf{U}\|, \|\mathbf{V}\| \leq O(1)$ and any deterministic vector $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{Td}$ such that $\|\mathbf{u}\|, \|\mathbf{v}\| \leq 1$, if one notes $\mathbf{B} = \frac{1}{Td} \mathbf{V}$ or $\mathbf{B} = \mathbf{u} \mathbf{v}^\top$, one can estimate:*

$$\left| \mathbb{E} \left[\text{tr}(\mathbf{B} \bar{\bar{\mathbf{Q}}}_\mathbf{U} \bar{\bar{\mathbf{Q}}}) \right] - \Psi(\mathbf{U}, \mathbf{B}) - \frac{1}{Td} \Psi(\mathbf{U})^\top \left(\mathbf{I}_n - \frac{1}{Td} \Psi \right)^{-1} \Psi(\mathbf{B}) \right| \leq O \left(\frac{1}{\sqrt{d}} \right) \quad (8)$$

where we noted:

- $\bar{\bar{\mathbf{Q}}} \equiv \bar{\bar{\mathbf{Q}}}_\delta(\mathbf{S}), \boldsymbol{\delta} = \boldsymbol{\delta}(\mathbf{S}),$
- $\Psi \equiv \frac{1}{Td} \left(\frac{\text{tr}(\mathbf{S}_i \bar{\bar{\mathbf{Q}}}_j \bar{\bar{\mathbf{Q}}})}{(1 + \delta_i)(1 + \delta_j)} \right)_{i,j \in [n]} \in \mathbb{R}^{n,n}$
- $\forall \mathbf{U} \in \mathbb{R}^{n \times n} : \Psi(\mathbf{U}) \equiv \frac{1}{Td} \left(\frac{\text{tr}(\mathbf{U} \bar{\bar{\mathbf{Q}}}_i \bar{\bar{\mathbf{Q}}})}{1 + \delta_i} \right)_{i \in [n]} \in \mathbb{R}^n$
- $\forall \mathbf{U}, \mathbf{V} \in \mathbb{R}^{n \times n} : \Psi(\mathbf{U}, \mathbf{V}) \equiv \frac{1}{Td} \text{tr} \left(\mathbf{U} \bar{\bar{\mathbf{Q}}}_\mathbf{V} \bar{\bar{\mathbf{Q}}}_\mathbf{U} \right) \in \mathbb{R}$

If there exist $T < n$ distinct matrices $\mathbf{C}_1, \dots, \mathbf{C}_T$ such that:

$$\{\mathbf{S}_1, \dots, \mathbf{S}_n\} = \{\mathbf{C}_1, \dots, \mathbf{C}_T\},$$

and if we denote $\forall t \in [T] n_t = \#\{i \in [n] \mid \mathbf{S}_i = \mathbf{C}_t\}$ and:

$$P \equiv \left(I_T - \left(\frac{n_t n_v}{(Td)^2} \frac{\text{tr}(\mathbf{S}_t \bar{\bar{\mathbf{Q}}}_s \bar{\bar{\mathbf{Q}}}_v)}{(1 + \delta_t)(1 + \delta_v)} \right)_{t,v \in [T]} \right)^{-1} \in \mathbb{R}^{T,T}$$

$$\forall \mathbf{U} \in \mathbb{R}^{Td \times Td} : \quad \bar{\bar{\mathbf{Q}}}_2(\mathbf{U}) \equiv \bar{\bar{\mathbf{Q}}}_\mathbf{U} \bar{\bar{\mathbf{Q}}}_\mathbf{U} + \frac{1}{(Td)^2} \sum_{t,v=1}^T \frac{\text{tr}(\mathbf{S}_t \bar{\bar{\mathbf{Q}}}_\mathbf{U} \bar{\bar{\mathbf{Q}}}_v) P_{t,v} \bar{\bar{\mathbf{Q}}}_s \bar{\bar{\mathbf{Q}}}_v \bar{\bar{\mathbf{Q}}}_t}{(1 + \delta_t)(1 + \delta_v)},$$

the result of Theorem 3 rewrites:

$$\left\| \mathbb{E} \left[\bar{\bar{\mathbf{Q}}}_\mathbf{U} \bar{\bar{\mathbf{Q}}}_\mathbf{U} \right] - \bar{\bar{\mathbf{Q}}}_2(\mathbf{U}) \right\| \leq O \left(\frac{1}{\sqrt{d}} \right) \quad (9)$$

Proof. Given $i \in [n]$, let us note $\mathbf{Z}_{-i} = (\mathbf{z}_1, \dots, \mathbf{z}_{i-1}, 0, \mathbf{z}_{i+1}, \dots, \mathbf{z}_n)$ and $\tilde{\mathbf{Q}}_{-i} = (\frac{1}{Td}\mathbf{Z}_{-i}\mathbf{Z}_{-i}^\top + \mathbf{I}_{Td})^{-1}$, then we have the identity:

$$\tilde{\mathbf{Q}} - \tilde{\mathbf{Q}}_{-i} = \frac{1}{Td}\tilde{\mathbf{Q}}\mathbf{z}_i\mathbf{z}_i^\top\tilde{\mathbf{Q}}_{-i} \quad \text{and} \quad \tilde{\mathbf{Q}}\mathbf{z}_i = \frac{\tilde{\mathbf{Q}}_{-i}\mathbf{z}_i}{1 + \frac{1}{Td}\mathbf{z}_i^\top\tilde{\mathbf{Q}}_{-i}\mathbf{z}_i}. \quad (10)$$

Given $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{Td}$, such that $\|\mathbf{u}\|, \|\mathbf{v}\| \leq 1$, let us express:

$$\mathbb{E} \left[\frac{1}{Td} \mathbf{u}^\top (\tilde{\mathbf{Q}} - \tilde{\mathbf{Q}}) \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \left(\frac{\mathbf{S}_i}{1 + \delta_i} - \mathbf{z}_i \mathbf{z}_i^\top \right) \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] \quad (11)$$

$$(12)$$

First, given $i \in [n]$, let us estimate thanks to (10):

$$\mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] = \mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] - \frac{1}{Td} \mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right]$$

Hölder inequality combined with Lemma 2 allows us to bound:

$$\frac{1}{Td} \left| \mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] \right| \leq \frac{1}{Td} \mathbb{E} \left[\left| \mathbf{u}^\top \tilde{\mathbf{Q}} \mathbf{z}_i \right|^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\left| \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right|^2 \right]^{\frac{1}{2}} \leq O \left(\frac{1}{d} \right),$$

one can thus deduce:

$$\mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] = \mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] + O \left(\frac{1}{d} \right) = \mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}}_{-i} \mathbf{v} \right] + O \left(\frac{1}{d} \right). \quad (13)$$

Second, one can also estimate thanks to Lemma 10:

$$\mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] = \mathbb{E} \left[\frac{\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v}}{1 + \frac{1}{Td} \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i} \right] = \mathbb{E} \left[\frac{\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v}}{1 + \delta_i} \right] + O \left(\frac{1}{\sqrt{d}} \right),$$

again thanks to Hölder inequality combined with Lemma 2 that allow us to bound:

$$\begin{aligned} & \mathbb{E} \left[\left| \frac{\delta_i - \frac{1}{Td} \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i}{(1 + \delta_i) \left(1 + \frac{1}{Td} \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \right)} \right| \left| \mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right| \right] \\ & \leq \mathbb{E} \left[\left| \delta_i - \frac{1}{Td} \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \right|^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\left| \mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right|^2 \right]^{\frac{1}{2}} \leq O \left(\frac{1}{\sqrt{d}} \right), \end{aligned}$$

The independence between \mathbf{z}_i and $\tilde{\mathbf{Q}}_{-i}$ (and $\tilde{\mathbf{Q}}$) then allow us to deduce (again with formula (10)):

$$\mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] = \mathbb{E} \left[\frac{\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}}_{-i} \mathbf{v}}{1 + \delta_i} \right] + \frac{1}{Td} \mathbb{E} \left[\frac{\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{v}}{1 + \delta_i} \right] + O \left(\frac{1}{\sqrt{d}} \right). \quad (14)$$

Let us inject (13) and (14) in (11) to obtain (again with an application of Hölder inequality and Lemma 2 that we do not detail this time):

$$\begin{aligned} \mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}} \left(\frac{\mathbf{S}_i}{1 + \delta_i} - \mathbf{z}_i \mathbf{z}_i^\top \right) \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{v} \right] &= \frac{1}{Td} \mathbb{E} \left[\frac{\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \mathbf{z}_i \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{v}}{\left(1 + \frac{1}{n} \mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \right)^2} \right] + O \left(\frac{1}{\sqrt{d}} \right), \\ &= \frac{1}{Td} \frac{\mathbb{E} \left[\mathbf{u}^\top \tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}}_{-i} \mathbf{v} \right]}{(1 + \delta_i)^2} \text{tr} \left(\mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \right) + O \left(\frac{1}{\sqrt{d}} \right), \end{aligned}$$

Putting all the estimations together, one finally obtains:

$$\left| \mathbb{E} \left[\tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \right] - \mathbb{E} \left[\tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \right] - \frac{1}{(Td)^2} \sum_{i=1}^n \frac{\text{tr} \left(\mathbf{S}_i \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}} \right)}{(1 + \delta_i)^2} \mathbb{E} \left[\tilde{\mathbf{Q}}_{-i} \mathbf{S}_i \tilde{\mathbf{Q}}_{-i} \right] \right| \leq O \left(\frac{1}{\sqrt{d}} \right) \quad (15)$$

One then see that if we introduce for any $\mathbf{V} \in \mathbb{R}^{n \times n}$ the block matrices:

- $\theta = \frac{1}{Td} \left(\frac{\mathbb{E}[\text{tr}(\mathbf{S}_j \tilde{\mathbf{Q}} \mathbf{S}_i \tilde{\mathbf{Q}}^Y)]}{(1+\delta_i)(1+\delta_j)} \right)_{i,j \in [n]} \in \mathbb{R}^{n \times n}$
- $\theta(\mathbf{V}) = \frac{1}{Td} \left(\frac{\mathbb{E}[\text{tr}(\mathbf{V} \tilde{\mathbf{Q}} \mathbf{S}_i \tilde{\mathbf{Q}}^Y)]}{1+\delta_i} \right)_{i \in [n]} \in \mathbb{R}^n$,
- $\theta(\mathbf{U}, \mathbf{V}) = \frac{1}{Td} \mathbb{E} \left[\text{tr}(\mathbf{V} \tilde{\mathbf{Q}} \mathbf{U} \tilde{\mathbf{Q}}^Y) \right] \in \mathbb{R}$,

then, if $\|\mathbf{V}\| \leq O(1)$, multiplying (15) with \mathbf{V} and taking the trace leads to:

$$\theta(\mathbf{U}, \mathbf{V}) = \Psi(\mathbf{U}, \mathbf{V}) + \frac{1}{Td} \Psi(\mathbf{U})^\top \theta(\mathbf{V}) + O\left(\frac{1}{\sqrt{d}}\right), \quad (16)$$

Now, taking $\mathbf{U} = \frac{\mathbf{S}_1}{1+\delta_1}, \dots, \frac{\mathbf{S}_n}{1+\delta_n}$, one gets the vectorial equation:

$$\theta(\mathbf{V}) = \Psi(\mathbf{V}) + \frac{1}{Td} \Psi \theta(\mathbf{V}) + O\left(\frac{1}{\sqrt{d}}\right),$$

When $(I_{Td} - \frac{1}{Td} \Psi)$ is invertible, one gets $\theta(\mathbf{V}) = (I_{Td} - \frac{1}{Td} \Psi)^{-1} \Psi(\mathbf{V}) + O\left(\frac{1}{\sqrt{d}}\right)$, and combining with (16), one finally obtains:

$$\theta(\mathbf{U}, \mathbf{V}) = \Psi(\mathbf{U}, \mathbf{V}) + \frac{1}{Td} \Psi(\mathbf{U})^\top (I_{Td} - \frac{1}{Td} \Psi)^{-1} \Psi(\mathbf{V}) + O\left(\frac{1}{\sqrt{d}}\right).$$

□

C.4 Estimation of the deterministic equivalent of \mathbf{Q}^2

Theorem 4. *Under the setting of Theorem 3, one can estimate:*

$$\|\mathbb{E}[\mathbf{Q}^2] - \mathbf{I}_n + \mathcal{D}_v\| \leq O\left(\frac{1}{\sqrt{d}}\right), \quad (17)$$

with, $\forall i \in [n]$:

$$v_i \equiv \frac{1}{Td} \frac{\text{tr}(\mathbf{S}_i \tilde{\mathbf{Q}})}{(1+\delta_i)^2} + \frac{1}{Td} \frac{\text{tr}(\mathbf{S}_i \tilde{\mathbf{Q}}_2(\mathbf{I}_n))}{(1+\delta_i)^2}$$

Proof. The justifications are generally the same as in the proof of Theorem 3, we will thus allow ourselves to be quicker in this proof.

Using the definition of $\mathbf{Q} = \left(\frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} + \mathbf{I}_n \right)^{-1}$, we have that

$$\frac{\mathbf{Z}^\top \mathbf{Z}}{Td} \mathbf{Q} = \left(\frac{\mathbf{Z}^\top \mathbf{Z}}{Td} + \mathbf{I}_n - \mathbf{I}_n \right) \left(\frac{\mathbf{Z}^\top \mathbf{Z}}{Td} + \mathbf{I}_n \right)^{-1} = \mathbf{I}_n - \mathbf{Q} \quad (18)$$

and one can then let appear $\tilde{\mathbf{Q}}$ thanks to the relation:

$$\mathbf{Z} \mathbf{Q} = \tilde{\mathbf{Q}} \mathbf{Z}, \quad (19)$$

that finally gives us:

$$\mathbf{Q} = \mathbf{I}_n - \frac{1}{Td} \mathbf{Z}^\top \mathbf{Z} \mathbf{Q} = \mathbf{I}_n - \frac{1}{Td} \mathbf{Z}^\top \tilde{\mathbf{Q}} \mathbf{Z}$$

One can then express:

$$\begin{aligned} \mathbf{Q}^2 &= \mathbf{I}_n - \frac{2}{Td} \mathbf{Z}^\top \tilde{\mathbf{Q}} \mathbf{Z} + \frac{1}{(Td)^2} \mathbf{Z}^\top \tilde{\mathbf{Q}} \mathbf{Z} \mathbf{Z}^\top \tilde{\mathbf{Q}} \mathbf{Z} \\ &= \mathbf{I}_n - \frac{1}{Td} \mathbf{Z}^\top \tilde{\mathbf{Q}} \mathbf{Z} - \frac{1}{Td} \mathbf{Z}^\top \tilde{\mathbf{Q}}^2 \mathbf{Z}. \end{aligned}$$

Given $i, j \in [n], i \neq j$, let us first estimate (thanks to Hölder inequality and Lemma 2):

$$\frac{1}{Td} \mathbb{E} \left[\mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{z}_j \right] = \frac{1}{Td} \frac{\mathbb{E} \left[\mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i,j} \mathbf{z}_j \right]}{(1 + \delta_i)(1 + \delta_j)} + O \left(\frac{1}{\sqrt{d}} \right) \leq O \left(\frac{1}{\sqrt{d}} \right),$$

since $\mathbb{E}[z_i] = \mathbb{E}[z_j] = 0$. Now, we consider the case $j = i$ to get:

$$\frac{1}{Td} \mathbb{E} \left[\mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{z}_i \right] = \frac{1}{Td} \frac{\mathbb{E} \left[\mathbf{z}_i^\top \tilde{\mathbf{Q}}_{-i} \mathbf{z}_i \right]}{(1 + \delta_i)^2} + O \left(\frac{1}{\sqrt{d}} \right) = \frac{1}{Td} \frac{\text{tr}(\mathbf{S}_i \tilde{\mathbf{Q}})}{(1 + \delta_i)^2} + O \left(\frac{1}{\sqrt{d}} \right).$$

As before, we know that $\frac{1}{Td} \mathbb{E} \left[\mathbf{z}_i^\top \tilde{\mathbf{Q}} \mathbf{z}_j \right] \leq O \left(\frac{1}{\sqrt{d}} \right)$ if $i \neq j$. Considering $i \in [n]$, we thus are left to estimate:

$$\frac{1}{Td} \mathbb{E} \left[\mathbf{z}_i^\top \tilde{\mathbf{Q}}^2 \mathbf{z}_j \right] = \frac{1}{Td} \frac{\text{tr}(\mathbf{S}_i \tilde{\mathbf{Q}}_2(\mathbf{I}_n))}{(1 + \delta_i)^2} + O \left(\frac{1}{\sqrt{d}} \right)$$

□

D Risk Estimation (Proof of Theorem 1)

D.1 Test Risk

The expected value of the MSE of the test data $\mathbf{x} \in \mathbb{R}^{T \times Td}$ concatenating the feature vector of all the tasks with the corresponding response variable $\mathbf{y} \in \mathbb{R}^{T \times Tq}$ reads as

$$\begin{aligned} \mathcal{R}_{test}^\infty &= \frac{1}{T} \mathbb{E} [\|\mathbf{y} - g(\mathbf{x})\|_2^2] \\ &= \frac{1}{T} \mathbb{E} \left[\left\| \frac{\mathbf{x}^\top \mathbf{W}}{\sqrt{Td}} + \varepsilon - \frac{\mathbf{x}^\top \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Y}}{Td} \right\|_2^2 \right] \\ &= \frac{1}{T} \mathbb{E} \left[\left\| \frac{\mathbf{x}^\top \mathbf{W}}{\sqrt{Td}} + \varepsilon - \frac{\mathbf{x}^\top \mathbf{A} \mathbf{Z} \mathbf{Q} \left(\frac{\mathbf{Z}^\top \mathbf{W}}{\sqrt{Td}} + \varepsilon \right)}{Td} \right\|_2^2 \right] \\ &= \frac{1}{T} \mathbb{E} \left[\left\| \frac{\mathbf{x}^\top \mathbf{W}}{\sqrt{Td}} + \varepsilon - \frac{\mathbf{x}^\top \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{W}}{Td \sqrt{Td}} - \frac{\mathbf{x}^\top \mathbf{A} \mathbf{Z} \mathbf{Q} \varepsilon}{Td} \right\|_2^2 \right] \\ &= \frac{1}{T} \mathbb{E} \left[\frac{\text{tr}(\mathbf{W}^\top \Sigma \mathbf{W})}{Td} - \frac{2 \text{tr}(\mathbf{W}^\top \Sigma \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{W})}{(Td)^2} + \text{tr}(\varepsilon^\top \varepsilon) + \frac{\text{tr}(\mathbf{W}^\top \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{A} \Sigma \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{W})}{(Td)^3} + \right. \\ &\quad \left. \frac{\text{tr}(\varepsilon^\top \mathbf{Q} \mathbf{Z}^\top \mathbf{A} \Sigma \mathbf{A} \mathbf{Z} \mathbf{Q} \varepsilon)}{(Td)^2} \right] \\ &= \frac{1}{T} \mathbb{E} \left[\frac{\text{tr}(\mathbf{W}^\top \Sigma \mathbf{W})}{Td} - \frac{2 \text{tr}(\mathbf{W}^\top \Sigma \mathbf{A}^{\frac{1}{2}} (\mathbf{I}_{Td} - \tilde{\mathbf{Q}}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W})}{Td} + \right. \\ &\quad \left. \text{tr}(\varepsilon^\top \varepsilon) + \frac{\text{tr}(\mathbf{W}^\top \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{A} \Sigma \mathbf{A} \mathbf{Z} \mathbf{Q} \mathbf{Z}^\top \mathbf{W})}{(Td)^3} + \frac{\text{tr}(\varepsilon^\top \mathbf{Q} \mathbf{Z}^\top \mathbf{A} \Sigma \mathbf{A} \mathbf{Z} \mathbf{Q} \varepsilon)}{(Td)^2} \right] \\ &= \frac{\text{tr}(\mathbf{W}^\top \Sigma \mathbf{W})}{Td} - 2 \frac{\text{tr}(\mathbf{W}^\top \Sigma \mathbf{A}^{\frac{1}{2}} (\mathbf{I}_{Td} - \tilde{\mathbf{Q}}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W})}{Td} + \text{tr}(\varepsilon^\top \varepsilon) + \frac{\text{tr}(\mathbf{W}^\top \Sigma \mathbf{W})}{Td} \\ &\quad - 2 \frac{\text{tr}(\mathbf{W}^\top \Sigma \mathbf{A}^{\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{-\frac{1}{2}} \mathbf{W})}{Td} + \frac{\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2(\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W}}{Td} + \frac{1}{Td} \text{tr}(\Sigma_N \tilde{\mathbf{Q}}_2) + O \left(\frac{1}{\sqrt{d}} \right) \end{aligned}$$

The test risk can be further simplified as

$$\mathcal{R}_{test}^\infty = \text{tr}(\Sigma_N) + \frac{\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2(\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W}}{Td} + \frac{\text{tr}(\Sigma_N \tilde{\mathbf{Q}}_2)}{Td} + O \left(\frac{1}{\sqrt{d}} \right)$$

D.2 Train Risk

In this section, we derive the asymptotic risk for the training data.

Theorem 5 (Asymptotic training risk). *Assuming that the training data vectors $\mathbf{x}_i^{(t)}$ and the test data vectors $\mathbf{x}^{(t)}$ are concentrated random vectors, and given the growth rate assumption (Assumption 2), it follows that:*

$$\mathcal{R}_{train}^\infty \leftrightarrow \frac{1}{Tn} \text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \bar{\bar{\mathbf{Q}}_2} \mathbf{A}^{-1/2} \mathbf{W} \right) - \frac{1}{Tn} \text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \bar{\bar{\mathbf{Q}}_2} (\mathbf{I}_{Td}) \mathbf{A}^{-1/2} \mathbf{W} \right) + \frac{1}{Tn} \text{tr} (\boldsymbol{\Sigma}_N \bar{\mathbf{Q}}_2)$$

Proof. We aim in this setting of regression, to compute the asymptotic theoretical training risk given by:

$$\mathcal{R}_{train}^\infty = \frac{1}{Tn} \mathbb{E} \left[\left\| \mathbf{Y} - \frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} \mathbf{Q} \mathbf{Y} \right\|_2^2 \right]$$

Using the definition of $\mathbf{Q} = \left(\frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} + \mathbf{I}_{Td} \right)^{-1}$, we have that

$$\frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} \mathbf{Q} = \left(\frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} + \mathbf{I}_{Td} - \mathbf{I}_{Td} \right) \left(\frac{\mathbf{Z}^\top \mathbf{A} \mathbf{Z}}{Td} + \mathbf{I}_{Td} \right)^{-1} = \mathbf{I}_{Td} - \mathbf{Q}$$

Plugging back into the expression of the training risk then leads to

$$\mathcal{R}_{train}^\infty = \frac{1}{Tn} \mathbb{E} [\text{tr} (\mathbf{Y}^\top \mathbf{Q}^2 \mathbf{Y})]$$

Using the definition of the linear generative model and in particular $\mathbf{Y} = \frac{\mathbf{Z}^\top \mathbf{W}}{\sqrt{Td}} + \boldsymbol{\varepsilon}$, we get

$$\begin{aligned} \mathcal{R}_{train}^\infty &= \frac{1}{Tn} \mathbb{E} \left[\text{tr} \left(\frac{1}{\sqrt{Td}} \mathbf{Z}^\top \mathbf{W} + \boldsymbol{\varepsilon} \right)^\top \mathbf{Q}^2 \left(\frac{1}{\sqrt{Td}} \mathbf{Z}^\top \mathbf{W} + \boldsymbol{\varepsilon} \right) \right] \\ &= \frac{1}{Tn} \frac{1}{Td} \mathbb{E} [\text{tr} (\mathbf{W}^\top \mathbf{Z} \mathbf{Q}^2 \mathbf{Z}^\top \mathbf{W})] + \frac{1}{Tn} \mathbb{E} [\text{tr} (\boldsymbol{\varepsilon}^\top \mathbf{Q}^2 \boldsymbol{\varepsilon})] \end{aligned}$$

To simplify this expression, we will introduced the so-called ‘‘coresolvent’’ defined as:

$$\tilde{\mathbf{Q}} = \left(\frac{\mathbf{A}^{\frac{1}{2}} \mathbf{Z} \mathbf{Z}^\top \mathbf{A}^{\frac{1}{2}}}{Td} + \mathbf{I}_{Td} \right)^{-1},$$

Employing the elementary relation $\mathbf{A}^{\frac{1}{2}} \mathbf{Z} \mathbf{Q} = \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{Z}$, one obtains:

$$\frac{1}{Td} \mathbf{Z} \mathbf{Q}^2 \mathbf{Z}^\top = \frac{1}{Td} \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{Z} \mathbf{Z}^\top \tilde{\mathbf{Q}} = \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}^2 \frac{\mathbf{A}^{\frac{1}{2}} \mathbf{Z} \mathbf{Z}^\top \mathbf{A}^{\frac{1}{2}}}{Td} \mathbf{A}^{-\frac{1}{2}} = \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{-\frac{1}{2}} - \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}^2 \mathbf{A}^{-\frac{1}{2}},$$

Therefore we further get

$$\mathcal{R}_{train}^\infty = \frac{1}{Tn} \mathbb{E} \left[\text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \tilde{\mathbf{Q}} \mathbf{A}^{-1/2} \mathbf{W} \right) \right] - \frac{1}{Tn} \mathbb{E} \left[\text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \tilde{\mathbf{Q}}^2 \mathbf{A}^{-1/2} \mathbf{W} \right) \right] + \frac{1}{Tn} \mathbb{E} [\text{tr} (\boldsymbol{\varepsilon}^\top \mathbf{Q}^2 \boldsymbol{\varepsilon})]$$

Using deterministic equivalents in Lemma 1, the training risk then leads to

$$\mathcal{R}_{train}^\infty = \frac{1}{Tn} \text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \bar{\bar{\mathbf{Q}}_2} \mathbf{A}^{-1/2} \mathbf{W} \right) - \frac{1}{Tn} \text{tr} \left(\mathbf{W}^\top \mathbf{A}^{-1/2} \bar{\bar{\mathbf{Q}}_2} (\mathbf{I}_{Td}) \mathbf{A}^{-1/2} \mathbf{W} \right) + \frac{1}{Tn} \text{tr} (\boldsymbol{\Sigma}_N \bar{\mathbf{Q}}_2) + O \left(\frac{1}{\sqrt{d}} \right)$$

□

E Interpretation and insights of the theoretical analysis

E.1 Analysis of the test risk

We recall the test risk as

$$\mathcal{R}_{test}^\infty = \text{tr} (\boldsymbol{\Sigma}_N) + \frac{\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \bar{\bar{\mathbf{Q}}_2} (\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W}}{Td} + \frac{\text{tr} (\boldsymbol{\Sigma}_N \bar{\mathbf{Q}}_2)}{Td} + O \left(\frac{1}{\sqrt{d}} \right)$$

The test risk is composed of a signal term of a signal term $\mathcal{S} = \frac{\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \bar{\bar{\mathbf{Q}}_2} (\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W}}{Td}$ and a noise term $\mathcal{N} = \frac{\text{tr} (\boldsymbol{\Sigma}_N \bar{\mathbf{Q}}_2)}{Td}$.

E.2 Interpretation of the signal term

Let's denote by $\bar{\Sigma} = \sum_{t=1}^T \frac{n_t d_t}{Td(1+\delta_t)^2} \Sigma^{(t)}$ and $\tilde{\Sigma} = \sum_{t=1}^T \frac{c_0}{1+\delta_t} \Sigma^{(t)}$. The signal term reads as

$$\mathcal{S} = \mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2(\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} \mathbf{W}.$$

Using the following identity,

$$\begin{aligned} \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2(\mathbf{A}) \mathbf{A}^{-\frac{1}{2}} &= \mathbf{A}^{-\frac{1}{2}} \bar{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} (\mathbf{I} + \bar{\Sigma}) \mathbf{A}^{\frac{1}{2}} \bar{\mathbf{Q}} \mathbf{A}^{-\frac{1}{2}} \\ &= (\mathbf{A} \tilde{\Sigma} + \mathbf{I})^{-1} (\mathbf{I} + \bar{\Sigma}) (\mathbf{A} \tilde{\Sigma} + \mathbf{I})^{-1} \end{aligned}$$

This finally leads to

$$\mathcal{S} = \mathbf{W}^\top (\mathbf{A} \tilde{\Sigma} + \mathbf{I})^{-1} (\mathbf{I} + \bar{\Sigma}) (\mathbf{A} \tilde{\Sigma} + \mathbf{I})^{-1} \mathbf{W}$$

The matrix $\mathcal{H} = (\mathbf{A} \tilde{\Sigma} + \mathbf{I})^{-1} (\mathbf{I} + \bar{\Sigma}) (\mathbf{A} \tilde{\Sigma} + \mathbf{I})^{-1}$ is responsible to amplifying the signal $\mathbf{W}^\top \mathbf{W}$ in order to let the test risk to decrease more or less. It is decreasing as function of the number of samples in the tasks n_t . Furthermore it is composed of two terms (from the independent training $\mathbf{W}_t^\top \mathbf{W}$) and the cross term $\mathbf{W}_t^\top \mathbf{W}_v$ for $t \neq v$. Both terms decreases as function of the number of samples n_t , smaller values of γ_t and increasing value of λ . The cross term depends on the matrix $\Sigma_t^{-1} \Sigma_v$ which materializes the covariate shift between the tasks. More specifically, if the features are aligned $\Sigma_t^{-1} \Sigma_v = \mathbf{I}$ and the cross term is maximal while for bigger Fisher distance between the covariance of the tasks, the correlation is not favorable for multi task learning. To be more specific the off-diagonal term of \mathcal{H} are responsible for the cross term therefore for the multi tasks and the diagonal elements are responsible for the independent terms.

To analyze more the element of \mathcal{H} , let's consider the case where $\Sigma^{(t)} = \mathbf{I}$ and $\gamma_t = \gamma$. In this case the diagonal and non diagonal elements \mathbf{D}_{IL} and \mathbf{C}_{MTL} are respectively given by

$$\mathbf{D}_{IL} = \frac{(c_0(\lambda + \gamma) + 1)^2 + c_0^2 \lambda^2}{(c_0(\lambda + \gamma) + 1)^2 - c_0^2 \lambda^2}, \quad \mathbf{C}_{MTL} = \frac{-2c_0 \lambda (c_0(\lambda + \gamma) + 1)}{(c_0(\lambda + \gamma) + 1)^2 - c_0^2 \lambda^2}$$

Both function are decreasing function of λ , $1/\gamma$ and c_0 .

E.3 Interpretation and insights of the noise terms

We recall the definition of the noise term \mathcal{N} as

$$\mathcal{N} = \text{tr} \left(\Sigma_N (\mathbf{A}^{-1} + \Sigma)^{-1} \right)$$

Now at the difference of the signal term there are no cross terms due to the independence between the noise of the different tasks. In this case on the diagonal elements of $(\mathbf{A}^{-1} + \Sigma)^{-1}$ matters. This diagonal term is increasing for an increasing value of the sample size, the value of λ . Therefore this term is responsible for the negative transfer. In the specific case where $\Sigma^{(t)} = \mathbf{I}_d$ and $\gamma_t = \gamma$ for all task t , the diagonal terms read as

$$\mathbf{N}_{NT} = \frac{(c_0(\lambda + \gamma)^2 + (\lambda + \gamma) - c_0 \lambda^2)^2 + \lambda^2}{((c_0(\lambda + \gamma) + 1)^2 - c_0^2 \lambda^2)^2}$$

E.4 Optimal Lambda

The test risk in the particular of identity covariance matrix can be rewritten as

$$\mathcal{R}_{test}^\infty = \mathbf{D}_{IL} (\|\mathbf{W}_1\|_2^2 + \|\mathbf{W}_2\|_2^2) + \mathbf{C}_{MTL} \mathbf{W}_1^\top \mathbf{W}_2 + \mathbf{N}_{NT} \text{tr} \Sigma_n.$$

Deriving $\mathcal{R}_{test}^\infty$ with respect to λ leads after some algebraic calculus to

$$\lambda^* = \frac{n}{d} SNR - \frac{\gamma}{2}$$

where the signal noise ratio is composed of the independent signal to noise ratio and the cross signal to noise ratio $SNR = \frac{\|\mathbf{W}_1\|_2^2 + \|\mathbf{W}_2\|_2^2}{\text{tr} \Sigma_n} + \frac{\mathbf{W}_1^\top \mathbf{W}_2}{\text{tr} \Sigma_n}$.

F Theoretical Estimations

F.1 Estimation of the training and test risk

The different theorems depends on the ground truth \mathbf{W} that needs to be estimated through $\hat{\mathbf{W}}$.

To estimate the test risk, one needs to estimate functionals of the form $\mathbf{W}^\top \mathbf{M} \hat{\mathbf{W}}$ and $\boldsymbol{\varepsilon}^\top \mathbf{M} \boldsymbol{\varepsilon}$ for any matrix \mathbf{M} . Using the expression of $\mathbf{W} = \mathbf{AZQY}$, we start computing $\hat{\mathbf{W}}^\top \mathbf{M} \hat{\mathbf{W}}$

$$\hat{\mathbf{W}}^\top \mathbf{M} \hat{\mathbf{W}} = \mathbf{Y}^\top \mathbf{QZ}^\top \mathbf{A} \mathbf{M} \mathbf{AZQY}$$

Using the generative model for $\mathbf{Y} = \frac{\mathbf{Z}^\top \mathbf{W}}{\sqrt{Td}} + \boldsymbol{\varepsilon}$, we obtain

$$\begin{aligned} \mathbb{E} \left[\hat{\mathbf{W}}^\top \mathbf{M} \hat{\mathbf{W}} \right] &= \mathbb{E} \left[\left(\frac{\mathbf{Z}^\top \mathbf{W}}{\sqrt{Td}} + \boldsymbol{\varepsilon} \right)^\top \mathbf{QZ}^\top \mathbf{A} \mathbf{M} \mathbf{AZQ} \left(\frac{\mathbf{Z}^\top \mathbf{W}}{\sqrt{Td}} + \boldsymbol{\varepsilon} \right) \right] \\ &= \frac{1}{Td} \mathbb{E} \left[\mathbf{W}^\top \mathbf{ZQZ}^\top \mathbf{A} \mathbf{M} \mathbf{AZQZ}^\top \mathbf{W} \right] + \mathbb{E} \left[\boldsymbol{\varepsilon}^\top \mathbf{QZ}^\top \mathbf{A} \mathbf{M} \mathbf{AZQ} \boldsymbol{\varepsilon} \right] \end{aligned}$$

Employing the elementary relation $\mathbf{A}^{\frac{1}{2}} \mathbf{ZQ} = \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{Z}$, one obtains:

$$\begin{aligned} \mathbb{E} \left[\hat{\mathbf{W}}^\top \mathbf{M} \hat{\mathbf{W}} \right] &= \frac{1}{Td} \mathbb{E} \left[\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{Z}^\top \mathbf{Z} \mathbf{A}^{\frac{1}{2}} \mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{Z} \mathbf{W} \right] + \mathbb{E} \left[\boldsymbol{\varepsilon}^\top \mathbf{QZ}^\top \mathbf{A} \mathbf{M} \mathbf{AZQ} \boldsymbol{\varepsilon} \right] \\ &= \mathbb{E} \left[\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \left(\mathbf{I} - \tilde{\mathbf{Q}} \right) \mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \left(\mathbf{I} - \tilde{\mathbf{Q}} \right) \mathbf{A}^{-\frac{1}{2}} \mathbf{W} \right] + \mathbb{E} \left[\boldsymbol{\varepsilon}^\top \mathbf{QZ}^\top \mathbf{A} \mathbf{M} \mathbf{AZQ} \boldsymbol{\varepsilon} \right] \\ &= \mathbb{E} \left[\mathbf{W}^\top \mathbf{M} \mathbf{W} \right] - 2 \mathbb{E} \left[\mathbf{W}^\top \mathbf{M} \mathbf{A}^{\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{-\frac{1}{2}} \mathbf{W} \right] + \mathbb{E} \left[\mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{-\frac{1}{2}} \mathbf{W} \right] \\ &\quad + \mathbb{E} \left[\boldsymbol{\varepsilon}^\top \mathbf{QZ}^\top \mathbf{A} \mathbf{M} \mathbf{AZQ} \boldsymbol{\varepsilon} \right] \end{aligned}$$

Using the deterministic equivalent of Lemma 1, we obtain

$$\begin{aligned} \hat{\mathbf{W}}^\top \mathbf{M} \hat{\mathbf{W}} &\leftrightarrow \mathbf{W}^\top \mathbf{M} \mathbf{W} - 2 \mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{W} + \text{tr} \Sigma_n \mathbf{M} \left(\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \right) + \mathbf{W}^\top \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2 \left(\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \right) \mathbf{A}^{-\frac{1}{2}} \mathbf{W} \\ &\leftrightarrow \mathbf{W}^\top \left(\mathbf{M} - 2 \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{M} + \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2 \left(\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \right) \mathbf{A}^{-\frac{1}{2}} \right) \mathbf{W} + \text{tr} \Sigma_n \tilde{\mathbf{Q}}_2 \left(\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \right) \\ &\leftrightarrow \mathbf{W}^\top \kappa(\mathbf{M}) \mathbf{W} + \text{tr} \Sigma_n \tilde{\mathbf{Q}}_2 \left(\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \right) \end{aligned}$$

where We define the mapping $\kappa : \mathbb{R}^{Td \times Td} \rightarrow \mathbb{R}^{q \times q}$ as follows

$$\kappa(\mathbf{M}) = \mathbf{M} - 2 \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{A}^{\frac{1}{2}} \mathbf{M} + \mathbf{A}^{-\frac{1}{2}} \tilde{\mathbf{Q}}_2 \left(\mathbf{A}^{\frac{1}{2}} \mathbf{M} \mathbf{A}^{\frac{1}{2}} \right) \mathbf{A}^{-\frac{1}{2}}.$$

F.2 Estimation of the noise covariance

The estimation of the noise covariance remains a technical challenge in this process. However, when the noise covariance is isotropic, it is sufficient to estimate only the noise variance. By observing that

$$\lim_{\lambda \rightarrow 0, \gamma \rightarrow \infty} \mathcal{R}_{train}^\infty = \sigma^2 \frac{\text{tr} \mathbf{Q}_2}{kn},$$

we can estimate the noise level from the training risk evaluated at large γ and $\lambda = 0$.

G Multivariate Time Series Forecasting

A version of this work with a stronger focus on time series applications can be found at [20].

G.1 Related Work

Multivariate Time Series Forecasting. Previous work has explored multivariate time series methods to enrich task representations and enhance diversity across time series applications [18].

Despite these advancements, random matrix theory, widely applied in multi-task learning, remains underutilized in multivariate time series contexts, particularly for multivariate time series forecasting (MTSF) [52, 54]. MTSF is common in applications like medical data [6], electricity consumption [49], temperatures [27], and stock prices [44]. Various methods, from classical tools [7, 45] and statistical approaches like ARIMA [3, 4] to deep learning techniques [8, 19, 32, 38, 40, 51, 53, 55, 56], have been developed for this task. Some studies prefer univariate models for multivariate forecasting [32], while others introduce channel-wise attention mechanisms [19]. We aim to enhance a univariate model by integrating a regularization technique from random matrix theory and multi-task learning.

G.2 Architecture and Training Parameters

Architectures without MTL regularization. We follow Chen et al. [8], Nie et al. [32], and to ensure a fair comparison of baselines, we apply the reversible instance normalization (RevIN) of Kim et al. [21]. All the baselines presented here are univariate i.e. no operation is performed by the network along the channel dimension. For the PatchTST baseline [32], we used the official implementation than can be found on [Github](#). The network used in Transformer follows the one in [19], using a single layer Transformer with one head of attention, while RevIN normalization and denormalization are applied respectively before and after the neural network function. The dimension of the model is $d_m = 16$ and remains the same in all our experiments. DLinearU is a single linear layer applied for each channel to directly project the subsequence of historical length into the forecasted subsequence of prediction length. It is the univariate extension of the multivariate DLinear, DLinearM, used in Zeng et al. [53]. The implementation of SAMformer can be found [here](#), and for the iTransformer architecture [here](#). These two multivariate models serve as baseline comparisons. We reported the results found in [19] and [24]. For all of our experiments, we train our baselines PatchTST, DLinearU and Transformer with the Adam optimizer [22], a batch size of 32 for the ETT datasets and 256 for the Weather dataset, and the learning rates summarized in Table 2.

Architectures with MTL Regularization. We implemented the univariate PatchTST, DLinearU, and Transformer baselines with MTL regularization. Initially, we scale the inputs twice using RevIN normalization. The first scaling is applied to the univariate components, and the second scaling is applied to the multivariate components. For each channel, we then apply our model without MTL regularization. The outputs are concatenated along the channel dimension, and this concatenation is flattened to form a matrix of shape (batch size, $q \times T$), where q is the prediction horizon and T is the number of channels. We then learn a square matrix W of shape $(q \times T) \times (q \times T)$ for projection and reshape the result to obtain an output of shape (batch size, q, T). This method can be applied on top of any univariate model. Our regularized loss has been introduced in 5.4.

Training parameters. The training/validation/test split is 12/4/4 months on the ETT datasets and 70%/20%/10% on the Weather dataset. We use a look-back window $d = 336$ for PatchTST and $d = 512$ for DLinearU and Transformer, using a sliding window with stride 1 to create the sequences. The training loss is the MSE. Training is performed during 100 epochs and we use early stopping with a patience of 5 epochs. For each dataset, baselines, and prediction horizon $H \in \{96, 192, 336, 720\}$, each experiment is run 3 times with different seeds, and we display the average of the test MSE over the 3 trials in Table 1.

Table 2: Learning rates used in our experiments.

Dataset	ETTh1/ETTh2	Weather
Learning rate	0.001	0.0001

G.3 Datasets

We conduct our experiments on 3 publicly available datasets of real-world time series, widely used for multivariate long-term forecasting [8, 32, 51]. The 2 Electricity Transformer Temperature datasets ETTh1, and ETTh2 [55] contain the time series collected by electricity transformers from July 2016 to July 2018. Whenever possible, we refer to this set of 2 datasets as ETT. Weather [27] contains the time series of meteorological information recorded by 21 weather indicators in 2020. It should be noted Weather is large-scale datasets. The ETT datasets can be downloaded [here](#) while the Weather

dataset can be downloaded [here](#). Table 3 sums up the characteristics of the datasets used in our experiments.

Table 3: Characteristics of the multivariate time series datasets used in our experiments.

Dataset	ETTh1/ETTh2	Weather
# features	7	21
# time steps	17420	52696
Granularity	1 hour	10 minutes

G.4 Additional Experiments.

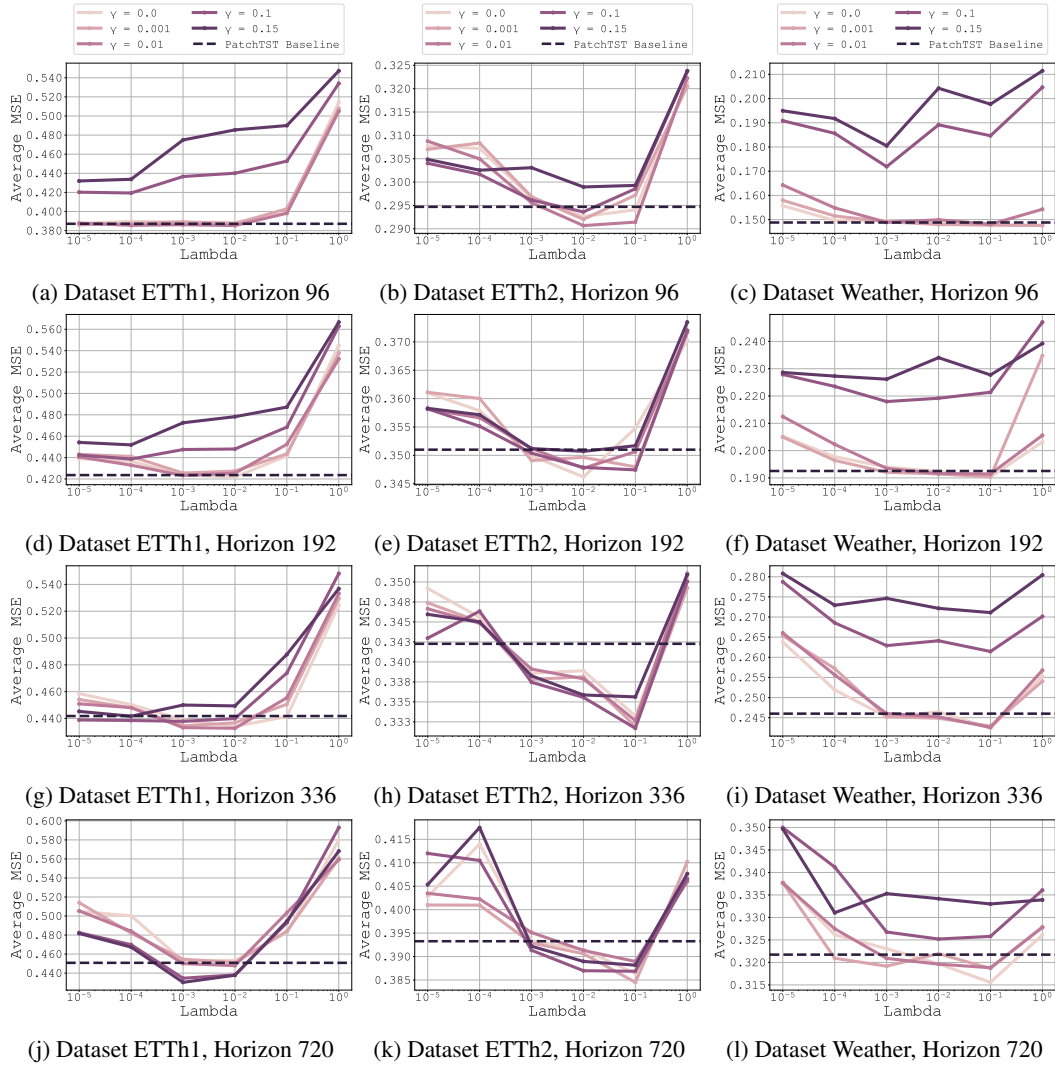


Figure 4: Results of our optimization method on different datasets and horizons averaged across 3 different seeds for each gamma and lambda values for the PatchTST baseline




Figure 5: Results of our optimization method on different datasets and horizons averaged across 3 different seeds for each gamma and lambda values for the DLinearU baseline

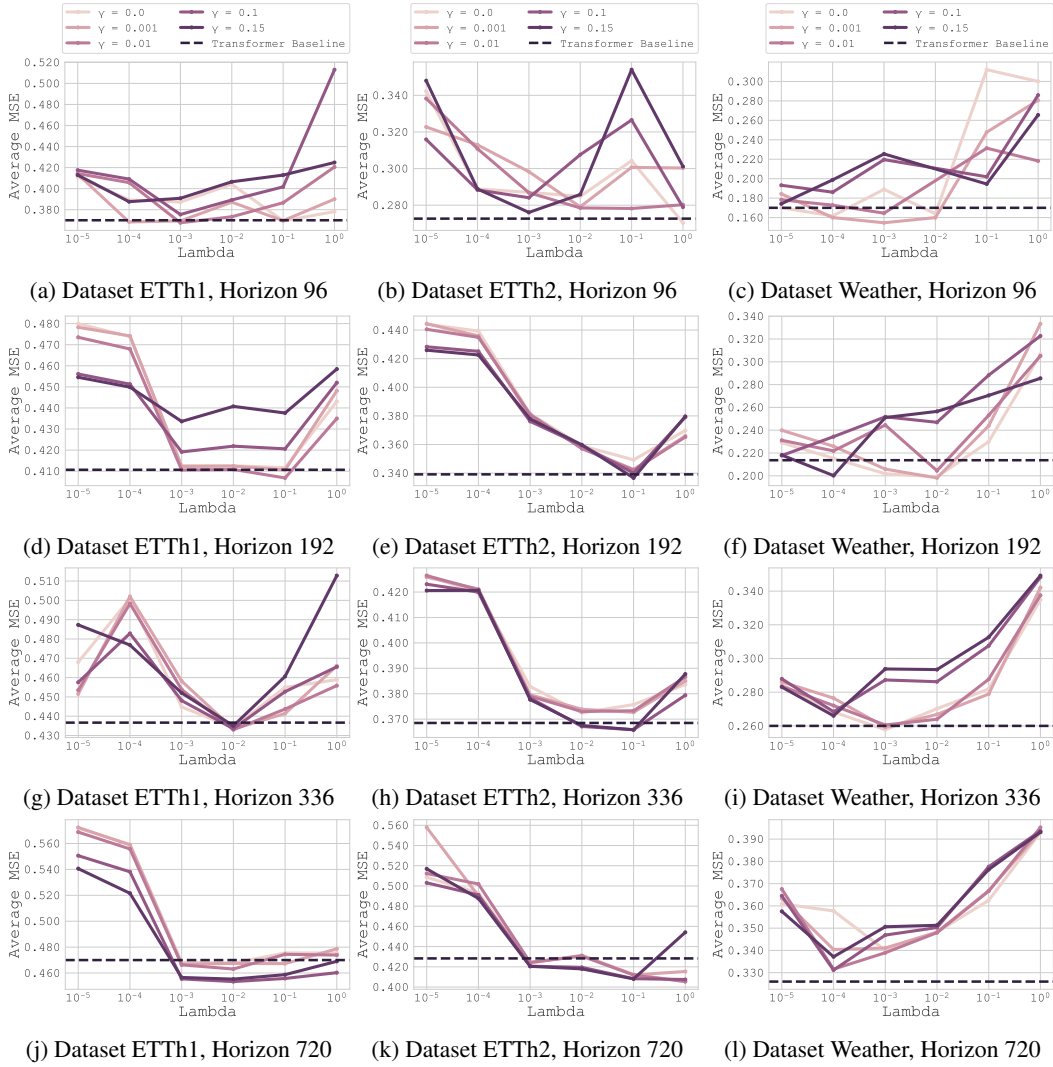


Figure 6: Results of our optimization method on different datasets and horizons averaged across 3 different seeds for each gamma and lambda values for the Transformer baseline

H Limitations

While the study provides valuable insights through its theoretical analysis within a linear framework, it is important to acknowledge its limitations. The linear approach serves as a solid foundation for understanding more complex models, but its practical applications may be constrained. Linear models, though mathematically tractable and often easier to interpret, might not fully capture the intricacies and nonlinear relationships present in real-world data, especially in the context of multivariate time series forecasting.

To address this limitation, we decided to extend our algorithm’s application to more complex models, specifically within the nonlinear setting of neural networks. This transition aims to evaluate whether the theoretical insights derived from the linear framework hold true empirically when applied to neural networks. As part of this endeavor, an optimal parameter λ was selected by an oracle, leading to promising results, as detailed in Section 5.4. This oracle-based selection underscores the potential efficacy of our approach when appropriately tuned, even in more complex, nonlinear contexts.

It is important to note that the limitations are not related to the real-world data itself, as our setting performs well in the context of multi-task regression for real-world data, as shown in Section 5.2. The difficulty arises from transitioning from a linear to a nonlinear model. The results in Section 5.4 are particularly encouraging, demonstrating that our method can improve upon univariate baselines by regularizing with an optimal λ , as indicated by our oracle. While the oracle provides an upper bound on performance, actual implementation would require robust methods for hyperparameter optimization in non-linear scenarios, which remains an open area for further research.

By expanding the scope of our theoretical framework to encompass nonlinear models, we pave the way for future work that could focus on the theoretical analysis of increasingly complex architectures

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The datasets used are open-source and all the implementation details are given to reproduce the experimental results. The code will be made available in case of acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No]

Justification: All details to reproduce experiments are given.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [No]

Justification: This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.