
Optimal Representation Size: High-Dimensional Analysis of Pretraining and Linear Probing

Anonymous Authors¹

Abstract

Learning reusable representations from unlabelled data is central to modern training pipelines, where large pretrained models are adapted to downstream tasks through fine-tuning or linear probing. We study this process in a tractable two-stage setting: representations are learned via principal component analysis on unlabelled data and reused for downstream linear regression on a separate labelled dataset. In the high-dimensional regime, we derive exact expressions for training and generalisation error as functions of representation dimensionality, unlabelled and labelled sample sizes, and task alignment. Our results show that pretrained representations strongly shape downstream generalisation and that the optimal representation size depends critically on the data regime: with abundant pretraining data but limited downstream supervision, maximally compressed representations are optimal, whereas with limited pretraining data, higher-dimensional representations generalise better. Beyond our idealised model, we observe similar phenomenology in autoencoders and pretrained LLMs. Altogether, our results identify when compression during pretraining improves downstream generalisation.

1. Introduction

Large language models, vision transformers, and related architectures have demonstrated language, reasoning and mathematical capabilities widely considered out of reach just a decade ago (OpenAI et al., 2024; Grattafiori et al., 2024; Wang et al., 2018; Hendrycks et al., 2021). This progress has been made possible by multi-stage training pipelines capable of leveraging massive amounts of unlabelled data: models are first pretrained on a large corpus

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

of unlabelled data, then fine-tuned on labelled tasks and finally refined, e.g., via reinforcement learning from human feedback (Christiano et al., 2017; Ouyang et al., 2022), low-rank adaptation (Hu et al., 2022) or test-time training (Sun et al., 2020; Liu et al., 2021; Akyürek et al., 2025). Such a paradigm unlocks remarkable sample efficiency, as large pretrained models generalise on new tasks where they have access only to thousands (or even hundreds) of samples (Aghajanyan et al., 2021; Howard & Ruder, 2018).

At their core, these training pipelines rely on a key fact: pretraining on large datasets yields rich representations that effectively transfer to downstream tasks. This idea is not new, as greedy layer-wise training and stacked autoencoders were designed to learn reusable representations in pretraining a decade before modern transformers (Bengio et al., 2006; Erhan et al., 2010; Vincent et al., 2010). However, different pretraining objectives emphasize different aspects of the data, and it remains unclear which features of a representation are most useful for downstream learning.

Characterizing these representations, how they support computation, transfer to new tasks and enable adaptive behaviour in artificial and biological systems, has been a longstanding focus in machine learning, as well as cognitive science and neuroscience (Bengio et al., 2013; Martin, 2007; Fusi et al., 2016). In machine learning, structured representations and prior knowledge can accelerate learning and enhance generalisation, but may also hinder adaptation, as exemplified by catastrophic interference in continual and reversal learning (Parisi et al., 2019; Kirkpatrick et al., 2017; Zenke et al., 2017; Braun et al., 2022). Similar trade-offs arise in biological systems, where representations must support flexible and efficient adaptation (Badre et al., 2021). A crucial property underlying these effects is the size (or dimensionality) of representations. Both high and low-dimensional representations have been observed across a range of brain regions and tasks (Boyle et al., 2024; Courellis et al., 2024; Mishchanchuk et al., 2024; Nogueira et al., 2023) as well as in machine learning models (Aghajanyan et al., 2021; Valeriani et al., 2023), yet a principled account of how dimensionality interacts with task and environment structure to shape learning performance and sample efficiency remains missing.

Here, we study a two-stage learning setting in which representations are first learned via principal component analysis (PCA) from unlabelled samples drawn from a spiked covariance model, and then reused by linear regression on an independent labelled dataset. The covariance structure plays a central role: it produces both a tractable model for analysis and controlled alignment between the pretraining and supervised training stages, thus enabling us to isolate how pretrained representations influence downstream generalisation. More precisely, our contributions are as follows:

- We derive exact asymptotic expressions for estimation, generalisation and training error (Theorem 4.1, 4.2 and 4.3, respectively). From this, we establish the improvement in generalisation coming from tuning the representation size, as well as the improvement due to adding more pretraining and finetuning data. For optimally-tuned errors, we identify regions where unlabelled samples are more valuable than labelled samples (Figure 2).
- We characterize when either (i) low-dimensional representations are useful for generalisation on downstream tasks, or (ii) tasks benefit from detailed high-dimensional representations (Figure 3). In particular, we identify phase transitions where compression becomes useful (Corollaries 4.4, 4.5), and describe their behaviour for an infinite number of pretraining samples (Corollary .13).
- We finally show that the theoretical predictions derived in our idealized setting extend to trained autoencoders when the number of samples is larger than the number of features, and further demonstrate that the insights unveiled by our theory persist in more realistic settings through linear probing of pretrained transformers (Figure 4).

2. Related work

High-dimensional regression. High-dimensional regression in proportional asymptotics has led to precise test-error formulae and explanations of benign overfitting and double descent (Dobriban & Wager, 2018; Trevor Hastie et al., 2022; Richards et al., 2020; Wu & Xu, 2020; Bartlett et al., 2020; Belkin et al., 2019; Advani et al., 2020). Principal component regression (PCR) restricts regression to a low-dimensional subspace, with performance governed by the number of retained components (Jolliffe, 1982; Xu & Hsu, 2019; Green & Romanov, 2025; Massy, 1965; Dhillon et al., 2013; Yao et al., 2015). In contrast, we study pretrained regression: PCs are estimated from an independent unlabelled sample and reused for a downstream task, changing both the asymptotic risk and the optimal representation size.

Pretraining and transfer learning. More broadly, the reuse of pretrained representations across tasks has been extensively studied in transfer learning, primarily in supervised settings where generalisation depends on task similarity or alignment with pretrained features (Pan & Yang, 2010; Zhuang et al., 2020; Gerace et al., 2022; Lampinen & Ganguli, 2019; Tahir et al., 2025; Yang et al., 2025; Song et al., 2024). These approaches typically do not model representations learned from unlabelled data. Complementary work establishes benefits of unsupervised and self-supervised pretraining (Lee et al., 2021; Azar & Nadler, 2024; Arora et al., 2019; Zhang & Hashimoto, 2021). Closest to our paper, recent work analyses PCA-based pretraining in spiked models (Jones-McCormick et al., 2025), but considers fixed representations without addressing the role of representation dimensionality in generalisation.

Representation dimensionality. A growing body of work highlights that pretrained representations often lie in low-dimensional subspaces, as evidenced by studies of intrinsic dimensionality (Aghajanyan et al., 2021; Valeriani et al., 2023), low-rank adaptation (Hu et al., 2022), and linear probing (Alain & Bengio, 2018). Recent theoretical work has begun to formalise aspects of this structure (Wang et al., 2026; Anguita et al., 2026). While these papers study the structure of pretrained features and how they support downstream prediction, they do not provide a closed-form characterisation of how the optimal retained dimension depends jointly on pretraining, downstream data, and task structure. Addressing this question through an analytically tractable theory is the focus of our work.

Spiked covariance models. The spiked covariance model (Johnstone, 2001) provides a tractable framework for studying low-dimensional structure in high-dimensional data. A key feature is the BBP phase transition (Jinho Baik et al., 2005; Marčenko & Pastur, 1967; Bai & Silverstein, 2010): above this threshold, empirical principal components align with the population spike, while below it they are asymptotically uninformative (Jinho Baik et al., 2005; Baik & Silverstein, 2006; Paul, 2007; Bai & Yao, 2012). This transition makes the spiked models well-suited for analysing how representation quality depends on sample size and signal strength. Accordingly, they have been used to analyse kernel methods, neural networks and feature learning (Ba et al., 2023; Ghorbani et al., 2020; Mousavi-Hosseini et al., 2023), as well as the emergence of low-rank structure during training (Ba et al., 2022; Cui et al., 2024; Dandi et al., 2025; Sonthalia et al., 2025). Finally, spiked covariances have been used to analyse PCR, showing how the projection onto empirical PCs has a regularising effect (Gedon et al., 2024). Our setting differs from these works: the signal subspace is estimated from an independent unlabelled pretraining sample and then used for a separate labelled task, allowing us to characterise the optimal representation dimension.




Figure 1. **Analytically tractable model of two-stage learning.** (a) In the pretraining stage, PCA extracts the top- m principal components from n_u unlabelled samples. The downstream task with n_l labelled samples is learned via regression on inputs projected through $\mathcal{P}_m = \mathbf{U}_m \mathbf{U}_m^\top$. Theoretically derived generalisation (b) and training errors (c) match numerical simulations.

3. Setup

We study two-stage learning with unlabelled pretraining followed by supervised learning:

1. In the pretraining stage, we observe inputs $\mathbf{X}_u \in \mathbb{R}^{n_u \times p}$ whose rows are i.i.d. $\mathcal{N}(\mathbf{0}, \Sigma)$, forming the *unlabelled* dataset \mathcal{D}_u . From this data, the model extracts a representation via PCA: it computes the top m principal components of \mathbf{X}_u , stacks them into $\mathbf{U}_m \in \mathbb{R}^{p \times m}$, and defines a projection of retained directions $\mathcal{P}_m = \mathbf{U}_m \mathbf{U}_m^\top$ with a discarded complement $\mathcal{P}_\perp = I - \mathcal{P}_m = \mathbf{U}_\perp \mathbf{U}_\perp^\top$. If $m > n_u$, the PCA subspace is completed with uniformly random orthonormal vectors.
2. In the downstream stage, we observe a *labelled* dataset $\mathcal{D}_l = \{\mathbf{X}_l, \mathbf{y}_l\}$, where $\mathbf{y}_l = \mathbf{X}_l \mathbf{w}^* + \xi \in \mathbb{R}^{n_l}$, the rows of $\mathbf{X}_l \in \mathbb{R}^{n_l \times p}$ are i.i.d. $\mathcal{N}(\mathbf{0}, \Sigma)$ and independent of \mathbf{X}_u , $\mathbf{w}^* \in \mathbb{R}^p$ is the ground-truth signal and $\xi \in \mathbb{R}^{n_l}$ is Gaussian noise with i.i.d. entries of variance σ^2 .

The pretrained regression model uses the representation learned from \mathcal{D}_u to fit the downstream task by regressing on the projected inputs (see Figure 1(a) for a schematic representation):

$$\hat{\mathbf{w}} = (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{y}_l = (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathbf{w}^* + (\mathbf{X}_l \mathcal{P}_m)^\dagger \xi, \quad (1)$$

where $(\cdot)^\dagger$ denotes the pseudoinverse. We denote by Π the projection onto the row space of $\mathbf{X}_l \mathcal{P}_m$, i.e., the subspace spanned by the projected training inputs. This setup follows (Green & Romanov, 2025), with the key difference that the projection subspace is learned from independent unlabelled data rather than the downstream dataset itself.

The quantities of interest are estimation $E^{\text{est}} = \mathbb{E}_\xi [\|\mathbf{w}^* - \hat{\mathbf{w}}\|^2]$, training $E^{\text{train}} = \frac{1}{n_l} \mathbb{E}_\xi [\|\mathbf{y}_l - \mathbf{X}_l \hat{\mathbf{w}}\|^2]$ and generalisation error $E^{\text{gen}} = \mathbb{E}_{\mathbf{x}, \mathbf{y}, \xi} [(y - \mathbf{x}^\top \hat{\mathbf{w}})^2]$, where the test point (\mathbf{x}, y) in E^{gen} is drawn from $\mathcal{N}(\mathbf{0}, \Sigma_{\text{new}})$.

We work in the high-dimensional regime where $p/n_l \rightarrow \gamma_l \in (0, \infty)$, $p/n_u \rightarrow \gamma_u \in (0, \infty)$, $m/p \rightarrow \alpha \in [0, 1]$ with α the fraction of dimensions to keep, and γ_u, γ_l the

aspect ratios for the datasets $\mathcal{D}_u, \mathcal{D}_l$. Due to the projection of the downstream inputs, we define the effective aspect ratio $\gamma_{\text{eff}} := \frac{m}{n_l} \rightarrow \alpha \gamma_l$ and assume $\gamma_{\text{eff}} \neq 1$.

The population covariance follows a rank-one spiked model $\Sigma = \mathbf{I}_p + (\lambda - 1) \mathbf{v} \mathbf{v}^\top$, where $\lambda > 1$ and the unit vector $\mathbf{v} \in \mathbb{R}^p$ are fixed. The alignment of \mathbf{w}^* with the spike eigenvector of Σ is captured by $(\mathbf{w}^{*\top} \mathbf{v})^2 / \|\mathbf{w}^*\|^2 \rightarrow \eta$, and the mass spectrum of \mathbf{w}^* on the eigenbasis of Σ is $\eta \delta_\lambda + (1 - \eta) \delta_1$. The test population covariance follows a spiked model with T spike directions $\{\mathbf{v}_{\text{new}, i}\}_{i=1}^T$ and eigenvalues $\{\nu_i\}_{i=1}^T$, i.e., $\Sigma_{\text{new}} = \mathbf{I}_p + \sum_{i=1}^T (\nu_i - 1) \mathbf{v}_{\text{new}, i} \mathbf{v}_{\text{new}, i}^\top$.

4. Theoretical analysis of two-stage learning

In this section, we derive asymptotic expressions for the estimation, generalisation, and training errors of the estimator in (1). We then characterise the phase transitions in the optimal representation dimensionality α , specifically, the regimes in which the optimum shifts from retaining all PCs to retaining only a few PCs, as a function of n_u and n_l . We further study how n_u , n_l , λ and the signal-to-noise ratio $\text{SNR} = \frac{\|\mathbf{w}^*\|^2}{\sigma^2}$ affect both the errors and the optimal α , as well as the limiting behaviour with infinite pretraining data ($\gamma_u = 0$). Proofs of all results are deferred to Appendix A.

4.1. Asymptotic expressions

The following deterministic limits characterise how the learned projection \mathcal{P}_m , its complement \mathcal{P}_\perp , and the downstream effective row-space projector Π capture fixed signal directions. Under the high-dimensional scaling assumptions of Section 3, for any deterministic vectors \mathbf{a}, \mathbf{b} with bounded normalized overlaps with the population spike directions, the bilinear projection terms satisfy

$$\frac{\mathbf{a}^\top \mathcal{P}_\perp \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} \xrightarrow{a.s.} \bar{p}_{\perp, \mathbf{a}, \mathbf{b}} \quad \frac{\mathbf{a}^\top \Pi \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} \xrightarrow{a.s.} \bar{\pi}_{\mathbf{a}, \mathbf{b}}, \quad (2)$$

with $\bar{p}_{\perp, \mathbf{a}, \mathbf{b}}$ and $\bar{\pi}_{\mathbf{a}, \mathbf{b}}$ given by the expressions in (36) and (43), involving integrals over limiting spectral distributions of $\mathbf{X}_u^\top \mathbf{X}_u / n_u$ and $(\mathbf{X}_l \mathcal{P}_m)^\top (\mathbf{X}_l \mathcal{P}_m) / n_l$ (see Lemmas .9 and .11 in Appendix A.6).

Theorem 4.1 (Asymptotic estimation error). As

165 $n_l, n_u, p, m \rightarrow \infty$ with $p/n_l \rightarrow \gamma_l$, $p/n_u \rightarrow \gamma_u$,
 166 $m/p \rightarrow \alpha$ and $\gamma_{\text{eff}} \neq 1$, the estimation error converges a.s.
 167 to

$$168 E_{\infty}^{\text{est}} = \bar{w}^* \left(\underbrace{1 - \bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*}}_{\text{missing signal}} + \underbrace{\frac{(\lambda - 1)^2}{\bar{\lambda}^2} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}^2 \bar{\pi}_{\mathbf{v}, \mathbf{v}}}_{\text{leaked signal}} \right) \quad (3)$$

$$171 + \underbrace{\frac{\min\{\gamma_{\text{eff}}, 1\}}{|\gamma_{\text{eff}} - 1|} (\bar{w}^* \bar{\sigma}_{\text{eff}}^2 + \sigma^2)}_{\text{variance}},$$

177 where \bar{w}^* is the deterministic limit of $\|\mathbf{w}^*\|^2$, $\bar{\lambda} = 1 +$
 178 $(\lambda - 1)(1 - \bar{p}_{\perp, \mathbf{v}, \mathbf{v}})$ is the effective spike eigenvalue, and
 179 the variance from the unrecoverable part of \mathbf{w}^* is $\bar{\sigma}_{\text{eff}}^2 =$
 180 $\bar{p}_{\perp, \mathbf{w}^*, \mathbf{w}^*} + \frac{\lambda - 1}{\lambda} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}^2$.

182 The expression in (3) admits a simple geometric interpre-
 183 tation: the estimator observes \mathbf{w}^* through the subspace in-
 184 duced by $\mathbf{X}_l \mathcal{P}_m$, and the error reflects how \mathbf{w}^* decomposes
 185 relative to this subspace. By decomposing into two compo-
 186 nents, aligned or orthogonal to \mathbf{U}_m , $\mathbf{w}^* = \mathbf{w}^*_{\parallel} + \mathbf{w}^*_{\perp}$, the
 187 error splits into three contributions: (a) the *missing signal*,
 188 (b) the *leaked signal*, and (c) the *variance* term.

- 191 (a) The term $1 - \bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*}$ captures the *missing signal*, i.e.
 192 directions that lie in the retained subspace but cannot
 193 be reliably estimated from finite samples.
- 195 (b) The key non-trivial effect is the *leak*. In an isotropic
 196 model, the components $\mathbf{X}_l \mathcal{P}_m$ and $\mathbf{X}_l \mathcal{P}_{\perp}$ are inde-
 197 pendent, so \mathbf{w}^*_{\perp} is completely invisible to the esti-
 198 mator. Here, the spike in the covariance couples
 199 these two parts of the data, so directions outside the
 200 retained subspace become correlated with it. As a re-
 201 sult, part of \mathbf{w}^*_{\perp} leaks into Π , where it appears as
 202 a spurious signal aligned with \mathbf{v} , producing the bias
 203 term $\frac{(\lambda - 1)^2}{\bar{\lambda}^2} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}^2 \bar{\pi}_{\mathbf{v}, \mathbf{v}}$. For a detailed discussion
 204 of parameter effects on the leak and a breakdown of
 205 different terms in the estimation error, see Figure 6 in
 206 Appendix B.
- 208 (c) The remaining part of \mathbf{w}^*_{\perp} does not couple back into
 209 the learned subspace. Together with label noise, this
 210 contributes to the *variance*, which exhibits a double
 211 descent (Belkin et al., 2019) controlled by $\gamma_{\text{eff}} = \alpha \gamma_l$.
 212 Thus, even without explicit regularisation in the re-
 213 gression step, the model mitigates double descent by
 214 controlling the representation size α , see Figure 6 in
 215 Appendix B.

217 **Theorem 4.2** (Asymptotic generalisation error). *Under the*
 218 *same regime, let $\rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}}$ be the deterministic limit of*

$\frac{\mathbf{w}^* \top \mathbf{v}_{\text{new}, i}}{\|\mathbf{w}^*\|}$. The generalisation error converges a.s. to

$$168 E_{\infty}^{\text{gen}} = E_{\infty}^{\text{est}} + \bar{w}^* \sum_{i=1}^T (\nu_i - 1) \left(\underbrace{\rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \bar{\pi}_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}}}_{\text{missing signal}} \right.$$

$$171 \left. - \underbrace{\frac{\lambda - 1}{\bar{\lambda}} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}} \bar{\pi}_{\mathbf{v}, \mathbf{v}_{\text{new}, i}}}_{\text{leaked signal}} \right)^2 + \sigma^2 \quad (4)$$

Thus, generalisation error equals estimation error plus the
 projection of $\mathbf{w}^* - \hat{\mathbf{w}}$ onto the spike directions of the test
 covariance. Each test spike contributes according to its
 strength $(\nu_i - 1)$, with the same missing- and leaked-signal
 decomposition as above. For additional discussions and a
 breakdown of different terms in the generalisation error, see
 Figure 7 in Appendix B.

Theorem 4.3 (Asymptotic training error). *Under the same*
regime, the training error converges a.s. to

$$E_{\infty}^{\text{train}} = (\bar{w}^* \bar{\sigma}_{\text{eff}}^2 + \sigma^2) \max(0, 1 - \gamma_{\text{eff}}).$$

This expression has the standard interpolation form. When
 $\gamma_{\text{eff}} > 1$, the model is overparameterised and perfectly
 fits the training data, yielding zero error. When $\gamma_{\text{eff}} < 1$,
 the model cannot interpolate, and the error decreases lin-
 early with γ_{eff} , with slope set by the total residual vari-
 ance $\bar{w}^* \bar{\sigma}_{\text{eff}}^2 + \sigma^2$.

Altogether, these results provide exact expressions for gener-
 alisation, estimation, and training errors, predicting several
 key trends. In Figure 1(b-c), we illustrate these errors as
 functions of the compression ratio α under varying data
 structures λ . The plots demonstrate an excellent agreement
 between theoretical predictions and experimental results
 for moderate values of n_u, n_l, p . Increasing λ reduces the
 bias from the missing signal, the leak and also the variance.
 However, it does not reduce the errors across α uniformly,
 as errors with low α benefit more from stronger structure
 than errors with high α . We validate our theory across a
 range of parameter settings, with experiments exploring vari-
 ations in the signal-to-noise ratio, the alignment between
 the spike \mathbf{v} and the signal \mathbf{w}^* , and the number of unlabelled
 and labelled samples (see Figure 5 in Appendix B).

Importantly, optimising the representation dimensionality
 reduces generalisation error across a broad region of param-
 eter space, both relative to low-dimensional representations
 with a single PC ($\alpha \approx 0$,¹ Figure 2(c)) and to full-rank rep-
 resentations ($\alpha = 1$, Figure 2(b)), with the optimum shown
 in Figure 2(a). These gains are particularly pronounced in

¹If m does not grow with p (e.g., $m = 1$), then $m/p \rightarrow \alpha = 0$.
 Thus, in simulations, $\alpha \approx 0$ refers to taking m not growing with p
 (e.g., retaining only the top principal component $m = 1$).




Figure 2. Benefits of optimal representation size. (a) Minimum achievable generalisation error with optimal α^* . (b) Gain in generalisation error from using α^* instead of all PCs ($\alpha = 1$). (c) Same gain relative to $\alpha \approx 0$. (d) Marginal rate of substitution between unlabelled and labelled data, $(\partial E_{\text{gen}}^{\text{gen}} / \partial n_u) / (\partial E_{\text{gen}}^{\text{gen}} / \partial n_l)$; values above 1 (red) indicate that unlabelled data reduces the error more than labelled data, and vice-versa below 1 (blue). (e) Generalisation-optimal representation size α^* . Simulation details are given in Appendix D.

regimes with substantial pretraining data and limited downstream samples, where standard PCR or regular regression fail to achieve comparable performance even in misaligned task settings (see Figure 12 in Appendix C.1 for a comparison).

Next, we study the trade-off between pretraining and supervision at the optimal α^* by quantifying how much error reduction is obtained from additional unlabelled versus labelled samples (Figure 2(d)). The resulting phase map reveals a nontrivial dependence on both data sources. In particular, there exists a broad interior region of the phase diagram—spanning moderate to large n_u and increasing n_l —in which additional pretraining data yields larger reductions in generalisation error than labelled samples. Interestingly, this regime is closely linked to the optimal representation size α^* that minimises generalisation error (see Figure 2(e)), which is discussed below.

4.2. Optimal representation size

We now study the optimal representation size quantified by the fraction of retained principal components, $\alpha = m/p$, that minimises generalisation and training error. We analyse how the resulting optimum depends on the interplay between unlabelled and labelled samples n_u and n_l , signal-to-noise ratio (SNR), spike strength λ , and task alignment η .

Generalisation error. We find that the optimal level of compression minimising generalisation error adapts to the availability of unlabelled and labelled data. For fixed SNR, λ and an aligned task $\mathbf{w}^* \parallel \mathbf{v}$, Figure 3(a) reveals the following three regimes.

- (i) Scarce n_l and n_u : not enough pretraining samples to reliably estimate PCs, thus downstream tasks should not be compressed ($\alpha^* \approx 1$).
- (ii) Large n_u , low n_l : maintain a compact representation subspace ($\alpha^* < 1$).
- (iii) High n_l : little compression needed ($\alpha^* \approx 1$), as the number of task samples is enough to learn.

We argue that regime (ii) is the most representative of modern machine-learning pipelines, where pretraining data is abundant but downstream task-specific data are limited. In practice, modern models are typically pretrained on large-scale datasets (Grattafiori et al., 2024; Biderman et al., 2023) and then fine-tuned on substantially smaller task-specific datasets (Wang et al., 2018). Altogether, we observe two phase transitions: one between regimes (i) and (ii), and another between regimes (ii) and (iii). We discuss them below in detail.

Phase transitions for generalisation error. We derive approximate analytical phase transition lines as a function of signal-to-noise ratio (SNR), spike strength (λ), and sample allocation between pretraining and downstream tasks (n_u and n_l). These approximations characterise the boundaries separating regimes in which different representation dimensions (α^*) are optimal, and they provide a tractable description of the underlying transitions observed in simulations.

Corollary 4.4 describes the transition between region (i) and (ii) when $\gamma_l > 1$. In this regime, the generalisation error exhibits a double descent, with two local minima: one at $\alpha = 1$, and one at small α (see Figure 2(b)). Although the exact location of the left minimum can in principle be obtained from the closed-form error expressions, the resulting formulas are not analytically transparent. Instead, Corollary A.7 provides an interpretable approximation by comparing the error at the two extremes $\alpha = 1$ and $\alpha = 0$, corresponding to retaining all components and only a fixed number of PCs.

Corollary 4.4 (Phase transition for $\gamma_l > 1$). *Let $\gamma_l > 1$ and assume $\Sigma_{\text{new}} = \Sigma$. Let $\text{SNR} \rightarrow S = \frac{\bar{w}^*}{\sigma^2}$ and define the overlap between the population and sample covariance spike eigenvector*

$$c(\gamma_u) = \mathbf{1}_{\gamma_u < (\lambda-1)^2} \frac{1 - \frac{\gamma_u}{(\lambda-1)^2}}{1 + \frac{\gamma_u}{\lambda-1}}.$$

Then, the transition between the small- α and large- α minima is approximated by the endpoint condition $E_{\text{gen}}^{\text{gen}}(0) =$




Figure 3. When is compression useful. (a) Heatmap of the generalisation-optimal α^* reveals distinct phases; white curves mark the theoretical phase-transition boundaries where low α is optimal. Phase boundaries in the (n_u, n_l) plane for varying SNR (b) varying λ (c) and varying η (d). Non-varying parameters are $\lambda = 5$, SNR = 9 and $\eta = 1$. Simulation details are given in Appendix D.

$E_\infty^{\text{gen}}(1)$, which yields a parametric curve

$$1 + \eta(\lambda - 1) - \eta \frac{\lambda^2 c(\gamma_u)}{1 + (\lambda - 1)c(\gamma_u)} = (1 - \eta) \frac{\gamma_l - 1}{\gamma_l} + \eta \frac{\lambda \gamma_l (\gamma_l - 1)}{(\gamma_l - 1 + \lambda)^2} + \frac{1}{S(\gamma_l - 1)}. \quad (5)$$

Along the curve (5), the two local minima of the generalisation error approximately coincide; crossing the curve then induces an abrupt transition in α^* . The derivation of this result is in Appendix A.7.

Corollary 4.5 characterises the transition between regions (ii) and (iii) in the underparametrised regime $\gamma_l < 1$. In this regime, the generalisation error exhibits a single minimum (see e.g. Figure 5(e,g,h) in Appendix B.1). The point $\alpha = 1$ is optimal as long as decreasing α increases the generalisation error, which is determined by the condition

$$\left. \frac{\partial E_\infty^{\text{gen}}}{\partial \alpha} \right|_{\alpha=1^-} = 0.$$

Corollary 4.5 (Phase transition for $\gamma_l < 1$). *Let $\gamma_l < 1$, SNR $\rightarrow S = \frac{\bar{w}^*}{\sigma^2}$ and denote by $D_{\mathbf{v}}(\gamma_u)$ the α -derivative of $\bar{p}_{\mathbf{v},\mathbf{v}}$ at $\alpha = 1$. Then, the solution $\alpha = 1$ is locally stable if and only if $\gamma_l \leq \gamma_l^{\text{crit}}$:*

$$\gamma_l^{\text{crit}} = \frac{S[(1 - \eta) + \eta D_{\mathbf{v}}(\gamma_u)]}{1 + S[(1 - \eta) + \eta D_{\mathbf{v}}(\gamma_u)]} \quad (6)$$

$$\text{where } D_{\mathbf{v}}(\gamma_u) = \begin{cases} \frac{\lambda \gamma_u}{(\lambda - 1 + \sqrt{\gamma_u})^2}, & \gamma_u < 1, \\ \frac{\gamma_u}{\gamma_u - 1 + \lambda}, & \gamma_u > 1. \end{cases}$$

Depending on the regime, the resulting boundary induces either a linear or nonlinear scaling of the critical number of labelled samples n_l with the unlabelled sample size n_u . Below this threshold, the optimal representation shifts to $\alpha < 1$, and compression improves downstream generalisation. The derivation of this result is in Appendix A.8.

Of particular interest is the limit $\gamma_u \rightarrow 0$, corresponding to infinite unlabelled data—a regime reminiscent of modern pretraining datasets (Grattafiori et al., 2024; OpenAI et al., 2024; Biderman et al., 2022). Here, errors and phase

boundaries are independent of γ_u and have simple dependencies on η , SNR, λ , and γ_l . Precise expressions are in Appendix A.9 (Corollary .12), where we also consider the limit of infinite labelled data, $\gamma_l \rightarrow 0$ (Corollary .13). In this last regime, the generalisation error simplifies (see Eq. (45)), the double-descent behaviour disappears, and the effective projector $(\mathbf{X}\mathcal{P}_m)^\dagger \mathbf{X}\mathcal{P}_m$ reduces to the population projector onto \mathcal{P}_m .

The full phase transition curve, derived in Corollaries 4.4 and 4.5, matches the phase transition observed in the heatmap of Figure 3(a) (shown in white; see also additional heatmaps in Figures 8 and 10 in Appendix B.2). Increasing the SNR progressively shrinks the region where compression is beneficial (Figure 3(b)), making full-rank representations ($\alpha = 1$) optimal over a wider range of sample sizes. The converse is true when increasing the spike strength λ or the task-spike alignment η (Figure 3(c,d)). These trends are evident from Corollary 4.5: when $\gamma_l < 1$, the stability boundary depends on $S(1 - \eta)$, so increasing SNR or decreasing η directly reduces γ_l^{crit} . Additional heatmaps for α in Appendix B.2 (Figures 8 and 10) further validate the joint dependence on SNR, λ and η , as well as the match with the theoretically derived transitions. In particular, increasing SNR, and decreasing λ and η all act to suppress the benefits of compression: for $\gamma_l < 1$ the phase boundary flattens, while for $\gamma_l > 1$ it shifts along the n_l axis, reducing the region where $\alpha < 1$ is optimal. Intuitively, a weaker spike is harder to estimate, lower alignment reduces the relevance of prior structure, and higher SNR makes the downstream task easy to learn from scratch. These effects diminish the value of compression, even when labelled data is abundant.

Training error. Our analysis further characterises the optimal value of α for minimising the training error. In the oversampled regime ($n_l > p$, corresponding to $\gamma_l < 1$), the optimal representation is full-rank, i.e., $\alpha = 1$. Conversely, in the undersampled regime ($n_l < p$, corresponding to $\gamma_l > 1$), the training error is equal to zero. In this case, the optimal dimensionality matches the sample size, i.e., $\alpha = n_l/p$, since increasing the dimension beyond this point provides no additional information. This is illustrated in

Figures 9 and 11 in Appendix B.2.

Altogether, our exact asymptotic analysis provides a theoretical framework to understand when low- or high-dimensional representations are optimal for downstream generalisation. We next validate such framework in practical settings.

5. Validation in autoencoders and pretrained transformers

In this section, we evaluate the extent to which theoretical predictions are valid beyond the idealised setting employed in their derivation. We first consider a controlled synthetic experiment in which the representation learned by an autoencoder approximates PCA. We then relax the assumptions of linearity and spiked isotropic covariance, allowing for a more general covariance structure. Finally, we show that our predictions extend to representations learned by modern neural networks.

5.1. Autoencoders as learned representations

We train the autoencoder on a pretraining task and use the learned representations for regression on downstream tasks. This approach is motivated by the fact that, in several regimes, autoencoders with a bottleneck dimension m recover the subspace spanned by the top- m principal components of the pretraining data (Refinetti & Goldt, 2023; Baldi & Hornik, 1989; Plaut, 2018; Bourlard & Kamp, 1988). We discuss below the regimes under which this correspondence holds.

Spiked isotropic covariance. We first train autoencoders on the spiked identity covariance structures considered in the theory. For linear autoencoders, the optimal bottleneck size, defined as the value minimising the generalisation error on a downstream task, closely follows our theoretical results (compare Figure 4(a) with Figure 3(a)). We highlight that the boundaries of the phase transitions in the heatmaps precisely match the predictions from Section 4.2. Nonlinear autoencoders exhibit the same qualitative behaviour, except when the number of pretraining samples is small, see Figure 4(b). This discrepancy is explained by a data-dependent transition: both linear and nonlinear autoencoders recover PCA only once sufficient pretraining data is available, see Figure 15 in Appendix C.2. In the low-sample regime, autoencoders instead tend to memorise training samples (Fumero et al., 2026; Jiang & Pehlevan, 2020; Radhakrishnan et al., 2020). Despite this, when $n_u < p$, higher-dimensional representations still yield lower errors on downstream tasks.

General covariance structure. For a spiked general covariance, the same qualitative behaviour persists. In fact, while phase boundaries shift, the overall structure is preserved: no compression when pretraining data is limited;

small bottlenecks for large pretraining and small downstream datasets; and intermediate-sized bottlenecks for large pretraining datasets. This is illustrated in Figure 4(c,d), and it is consistent with PCA-based regression under general covariance, see Figure 13 in Appendix C.2. Furthermore, when optimising for training error, the results closely match the spiked identity case (Figure 13 in Appendix C.2). At lower signal-to-noise ratios, the regime where compression is optimal expands (Figure 14 in Appendix C.2), in agreement with the theoretical analysis. Nonlinear autoencoders again exhibit deviations from PCA at low pretraining sample sizes.

Overall, our empirical results show that the theoretical predictions derived in the idealised setting extend to trained linear and nonlinear autoencoders when enough pretraining data is available, and that the qualitative phase structure persists under general covariance spectra.

5.2. Linear probing transformers

We next examine whether representations learned by modern architectures exhibit similar dimensionality effects under linear probing. We use models from the Pythia family (Biderman et al., 2023), which provide access to intermediate checkpoints throughout pretraining. As these models are trained for ~ 1.5 epochs on the Pile dataset (Gao et al., 2020; Biderman et al., 2022), the training step serves as a proxy for the number of pretraining samples n_u . The learned representations exhibit a spiked covariance structure, with a clear bulk and a small number of outlying eigenvalues (Figure 4(f)), and the spikes become more pronounced as training progresses, see Figure 18 in Appendix C.3. For each checkpoint, we project representations onto the top- m principal components and train linear probes on a downstream sentiment classification task (SST-2 from GLUE (Wang et al., 2018)), selecting m to optimise validation or training accuracy.

The dependence on dataset size is consistent with theoretical predictions, see Figure 4(e). For small downstream datasets, models that are pretrained longer benefit from stronger compression, see Figure 4(g). For larger downstream datasets, performance improves monotonically with m , indicating that little or no compression is optimal (see the right-most plot of Figure 17 in Appendix C.3). For extremely small downstream datasets (~ 100 samples), the optimal dimensionality is high again (see the left-most plot in Figure 17 in Appendix C.3), consistent with the weak alignment between the task signal and dominant pretraining directions. Finally, tuning the representation dimensionality via PCA consistently matches or outperforms using the full feature space, see Figure 19 in Appendix C.3.

In summary, these results show that real-world representations exhibit a theoretically predictable qualitative phase

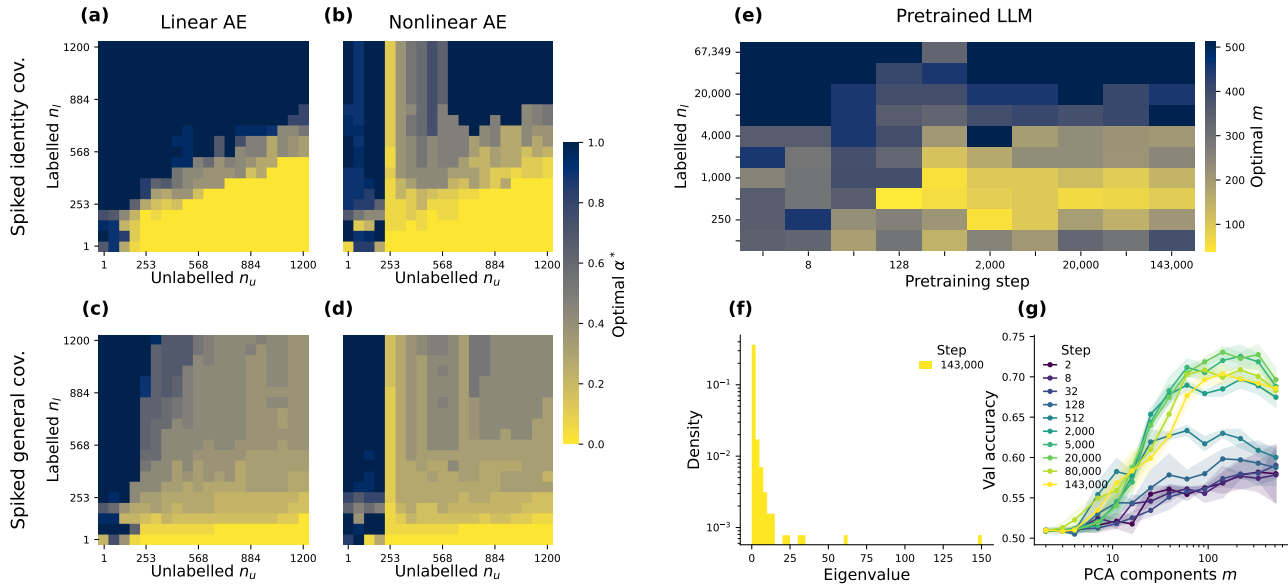


Figure 4. Optimal representation in autoencoders and pretrained LLMs. Optimal representation size as a function of unlabelled (n_u) and labelled (n_l) sample sizes. Linear autoencoders are considered in panels (a,c) and nonlinear ones in panels (b,d). Autoencoders are trained on Gaussian data with either spiked identity covariance $\Sigma = \mathbf{I}_p + \lambda \mathbf{v}\mathbf{v}^\top$ (panels (a,b)) or a spiked Toeplitz covariance $\Sigma = \mathbf{H} + \lambda \mathbf{v}\mathbf{v}^\top$, where $H_{i,j} = \rho^{|i-j|}$ and $\rho = 0.5$ (panels (b,d)). (e) In a language transformer (PYTHIA models), PCA on representations also improves downstream linear probing on the `sst2` task from the GLUE benchmark. (f) Eigenvalue spectrum of last-token representations from the final hidden layer, showing a bulk with a few spikes. (g) Validation accuracy versus representation size m for a small downstream dataset ($n_l = 2000$); colors denote different checkpoints and shaded regions indicate standard deviation across seeds. Simulation details are given in Appendix D.

structure when optimising dimensionality for downstream generalisation.

6. Discussion

We present a theoretical model quantifying how representation dimensionality affects training and generalisation error on downstream tasks. We demonstrate that (i) tuning this dimensionality significantly improves generalisation compared to full representations or naive fixed dimensionality reduction, and that (ii) the optimal dimensionality is determined by the amount of pretraining and downstream data, the task signal-to-noise ratio, the strength of structure in the inputs, and the alignment between task-relevant signals and dominant covariance directions. In particular, when pretraining data is abundant and downstream data is limited, the optimal representations are maximally compressed. Furthermore, we unveil a trade-off between pretraining and supervision, identifying regimes where unlabelled data yields larger gains than labelled samples. These results provide a principled explanation for why large models trained on extensive data and exhibiting low intrinsic dimensionality (Aghajanyan et al., 2021), are especially effective for fine-tuning and linear probing.

The question of optimal representations is central to machine learning as well as neuroscience, where internal representations support categorisation, reasoning, planning, naviga-

tion, and decision-making (Stokes et al., 2013; Genovesio et al., 2005; Shima & Tanji, 2000). A key property of these representations is their dimensionality. High-dimensional representations may support flexible readout across tasks (Raposo et al., 2014; Tang et al., 2019; Rigotti et al., 2013; Bernardi et al., 2020), whereas low-dimensional representations are associated with abstraction, efficient coding, or structured prior knowledge (Tye et al., 2024; Nogueira et al., 2023). Our results provide a quantitative characterization of how task structure and data availability determine the optimal representation tradeoff in a simplified setting, offering predictions for how biological systems might organize representations for effective learning.

Several limitations motivate future work: relaxing the shared-spike assumption to capture data mismatch, extending the analysis to ridge-regularised learning, and considering richer covariance models. Another important direction is to replace the closed-form regression estimator with gradient-based training, enabling the study of optimisation-induced bias and its dependence on initialisation from pre-training. Finally, extending the framework beyond PCA-based projections to learned feature representations would allow for a closer connection to empirical observations in deep networks.

Impact Statement

This paper presents theoretical work aimed at advancing the understanding of representation learning and transfer in machine learning systems. Potential positive impacts include improving the efficiency and sample efficiency of pretrained models by identifying when compressed representations improve downstream generalisation. We do not identify specific ethical concerns beyond those generally associated with advances in machine learning methods and representation learning.

References

- Advani, M. S., Saxe, A. M., and Sompolinsky, H. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132, 2020.
- Aghajanyan, A., Gupta, S., and Zettlemoyer, L. Intrinsic Dimensionality Explains the Effectiveness of Language Model Fine-Tuning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021.
- Akyürek, E., Damani, M., Zweiger, A., Qiu, L., Guo, H., Pari, J., Kim, Y., and Andreas, J. The surprising effectiveness of test-time training for few-shot learning. In *International Conference on Machine Learning*. PMLR, 2025.
- Alain, G. and Bengio, Y. Understanding intermediate layers using linear classifier probes. *arXiv preprint arXiv:1610.01644*, 2018.
- Anguita, N., Locatello, F., Saxe, A. M., Mondelli, M., Mancini, F., Lippl, S., and Domine, C. A theory of how pretraining shapes inductive bias in fine-tuning. *arXiv preprint arXiv:2602.20062*, 2026.
- Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- Azar, E. and Nadler, B. Semi-supervised sparse gaussian classification: Provable benefits of unlabeled data. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Ba, J., Erdogdu, M. A., Suzuki, T., Wang, Z., Wu, D., and Yang, G. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In *Advances in Neural Information Processing Systems*, volume 35, 2022.
- Ba, J., Erdogdu, M. A., Suzuki, T., Wang, Z., and Wu, D. Learning in the presence of low-dimensional structure: a spiked random matrix perspective. In *Advances in Neural Information Processing Systems*, volume 36, 2023.
- Badre, D., Bhandari, A., Keglovits, H., and Kikumoto, A. The dimensionality of neural representations for control. *Current Opinion in Behavioral Sciences*, 38, 2021.
- Bai, Z. and Silverstein, J. W. *Spectral Analysis of Large Dimensional Random Matrices*. Springer Series in Statistics. Springer New York, 2010.
- Bai, Z. and Yao, J. On sample eigenvalues in a generalized spiked population model. *Journal of Multivariate Analysis*, 106, 2012.
- Baik, J. and Silverstein, J. W. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of multivariate analysis*, 97, 2006.
- Baldi, P. and Hornik, K. Neural networks and principal component analysis: Learning from examples without local minima. *Neural Networks*, 2, 1989.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117, 2020.
- Belkin, M., Hsu, D., Ma, S., and Mandal, S. Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proceedings of the National Academy of Sciences*, 116, 2019.
- Bengio, Y., Lamblin, P., Popovici, D., and Larochelle, H. Greedy layer-wise training of deep networks. In *Advances in Neural Information Processing Systems*, volume 19, 2006.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2013.
- Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., and Salzman, C. D. The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell*, 183, 2020.
- Biderman, S., Bicheno, K., and Gao, L. Datasheet for the pile. *arXiv preprint arXiv:2201.07311*, 2022.
- Biderman, S., Schoelkopf, H., Anthony, Q. G., Bradley, H., O'Brien, K., Hallahan, E., Khan, M. A., Purohit, S., Prashanth, U. S., Raff, E., et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, 2023.
- Bourlard, H. and Kamp, Y. Auto-association by multilayer perceptrons and singular value decomposition. *Biological Cybernetics*, 59, 1988.

- 495 Boyle, L. M., Posani, L., Irfan, S., Siegelbaum, S. A., and
 496 Fusi, S. Tuned geometries of hippocampal representa-
 497 tions meet the computational demands of social memory.
 498 *Neuron*, 112, 2024.
- 499 Braun, L., Dominé, C. C. J., Fitzgerald, J. E., and Saxe,
 500 A. M. Exact learning dynamics of deep linear networks
 501 with prior knowledge. In *Advances in Neural Information*
 502 *Processing Systems*, 2022.
- 504 Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg,
 505 S., and Amodei, D. Deep reinforcement learning from
 506 human preferences. In *Advances in Neural Information*
 507 *Processing Systems*, volume 30, 2017.
- 509 Courellis, H. S., Minxha, J., Cardenas, A. R., Kimmel,
 510 D. L., Reed, C. M., Valiante, T. A., Salzman, C. D.,
 511 Mamelak, A. N., Fusi, S., and Rutishauser, U. Abstract
 512 representations emerge in human hippocampal neurons
 513 during inference. *Nature*, 632, 2024.
- 514 Cui, H., Pesce, L., Dandi, Y., Krzakala, F., Lu, Y. M., Zde-
 515 borová, L., and Loureiro, B. Asymptotics of feature
 516 learning in two-layer networks after one gradient-step. In
 517 *International Conference on Machine Learning*, 2024.
- 519 Dandi, Y., Pesce, L., Cui, H., Krzakala, F., Lu, Y., and
 520 Loureiro, B. A random matrix theory perspective on the
 521 spectrum of learned features and asymptotic generaliza-
 522 tion capabilities. In *International Conference on Artificial*
 523 *Intelligence and Statistics*, pp. 2224–2232. PMLR, 2025.
- 524 Dhillon, P. S., Foster, D. P., Kakade, S. M., and Ungar,
 525 L. H. A risk comparison of ordinary least squares vs
 526 ridge regression. *Journal of Machine Learning Research*,
 527 14, 2013.
- 529 Dobriban, E. and Wager, S. High-dimensional asymptotics
 530 of prediction: Ridge regression and classification. *The*
 531 *Annals of Statistics*, 2018.
- 533 Erhan, D., Bengio, Y., Courville, A., Manzagol, P.-A., Vin-
 534 cent, P., and Bengio, S. Why does unsupervised pre-
 535 training help deep learning? *Journal of Machine Learn-*
 536 *ing Research*, 11, 2010.
- 537 Fumero, M., Moschella, L., Rodolà, E., and Locatello, F.
 538 Navigating the latent space dynamics of neural models.
 539 In *The Fourteenth International Conference on Learning*
 540 *Representations*, 2026.
- 542 Fusi, S., Miller, E. K., and Rigotti, M. Why neurons mix:
 543 high dimensionality for higher cognition. *Current Opin-*
 544 *ion in Neurobiology*, 37, 2016.
- 546 Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T.,
 547 Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N.,
 548 et al. The Pile: An 800GB dataset of diverse text for
 549 language modeling. *arXiv preprint arXiv:2101.00027*,
 2020.
- Gedon, D., Ribeiro, A. H., and Schön, T. B. No double
 descent in principal component regression: A high-
 dimensional analysis. In *International Conference on*
Machine Learning, 2024.
- Genovesio, A., Brasted, P. J., Mitz, A. R., and Wise, S. P.
 Prefrontal cortex activity related to abstract response
 strategies. *Neuron*, 47, 2005.
- Gerace, F., Saglietti, L., Mannelli, S. S., Saxe, A., and
 Zdeborová, L. Probing transfer learning with a model of
 synthetic correlated datasets. *Machine Learning: Science*
and Technology, 2022.
- Ghorbani, B., Mei, S., Misiakiewicz, T., and Montanari, A.
 When do neural networks outperform kernel methods?
 In *Advances in Neural Information Processing Systems*,
 volume 33, 2020.
- Grattafiori, A. et al. The Llama 3 Herd of Models. *arXiv*
preprint arXiv:2407.21783, 2024.
- Green, A. and Romanov, E. The high-dimensional asymp-
 totics of principal component regression. *The Annals of*
Statistics, 53, 2025.
- Hendrycks, D., Burns, C., Kadavath, S., Arora, A., Basart,
 S., Tang, E., Song, D., and Steinhardt, J. Measuring
 mathematical problem solving with the math dataset. In
Proceedings of the Neural Information Processing Sys-
tems Track on Datasets and Benchmarks, volume 1, 2021.
- Howard, J. and Ruder, S. Universal language model fine-
 tuning for text classification. In *Proceedings of the 56th*
Annual Meeting of the Association for Computational
Linguistics (Volume 1: Long Papers), 2018.
- Hu, E. J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y.,
 Wang, S., Wang, L., and Chen, W. LoRA: Low-rank
 adaptation of large language models. In *International*
Conference on Learning Representations, 2022.
- Jiang, Y. and Pehlevan, C. Associative memory in iterated
 overparameterized sigmoid autoencoders. In *Proceed-*
ings of the 37th International Conference on Machine
Learning, 2020.
- Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase
 transition of the largest eigenvalue for nonnull complex
 sample covariance matrices. *The Annals of Probability*,
 33, 2005.
- Johnstone, I. M. On the distribution of the largest eigenvalue
 in principal components analysis. *The Annals of Statistics*,
 29(2), 2001.

- 550 Jolliffe, I. T. A Note on the Use of Principal Components in
551 Regression. *Applied Statistics*, 31, 1982.
- 552 Jones-McCormick, T., Jagannath, A., and Sen, S. Provable
553 benefits of unsupervised pre-training and transfer learning
554 via single-index models. In *Proceedings of the 42nd*
555 *International Conference on Machine Learning*, volume
556 267. PMLR, 2025.
- 557 Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Des-
558 jardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T.,
559 Grabska-Barwinska, A., et al. Overcoming catastrophic
560 forgetting in neural networks. *Proceedings of the national*
561 *academy of sciences*, 114, 2017.
- 562 Lampinen, A. K. and Ganguli, S. An analytic theory of
563 generalization dynamics and transfer learning in deep
564 linear networks. In *International Conference on Learning*
565 *Representations*, 2019.
- 566 Lee, J. D., Lei, Q., Saunshi, N., and Zhuo, J. Predicting
567 what you already know helps: Provable self-supervised
568 learning. In Beygelzimer, A., Dauphin, Y., Liang, P., and
569 Vaughan, J. W. (eds.), *Advances in Neural Information*
570 *Processing Systems*, 2021.
- 571 Liu, Y., Kothari, P., Van Delft, B., Bellot-Gurlet, B., Mordan,
572 T., and Alahi, A. Ttt++: When does self-supervised
573 test-time training fail or thrive? In *Advances in Neural*
574 *Information Processing Systems*, volume 34, 2021.
- 575 Martin, A. The representation of object concepts in the
576 brain. *Annu. Rev. Psychol.*, 58, 2007.
- 577 Marčenko, V. A. and Pastur, L. A. Distribution of eigenval-
578 ues for some sets of random matrices. *Mathematics of*
579 *the USSR-Sbornik*, 1, 1967.
- 580 Massy, W. F. Principal components regression in exploratory
581 statistical research. *Journal of the American Statistical*
582 *Association*, 60, 1965.
- 583 Mishchanchuk, K., Gregoriou, G., Qü, A., Kastler, A., Huys,
584 Q. J. M., Wilbrecht, L., and MacAskill, A. F. Hidden state
585 inference requires abstract contextual representations in
586 the ventral hippocampus. *Science*, 386, 2024.
- 587 Mousavi-Hosseini, A., Wu, D., Suzuki, T., and Erdogdu,
588 M. A. Gradient-based feature learning under structured
589 data. In *Advances in Neural Information Processing*
590 *Systems*, volume 36, 2023.
- 591 Nogueira, R., Rodgers, C. C., Bruno, R. M., and Fusi, S. The
592 geometry of cortical representations of touch in rodents.
593 *Nature Neuroscience*, 26, 2023.
- 594 OpenAI et al. GPT-4 Technical Report. *arXiv preprint*
595 *arXiv:2303.08774*, 2024.
- 596 Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.,
597 Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A.,
598 Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M.,
599 Askell, A., Welinder, P., Christiano, P. F., Leike, J., and
600 Lowe, R. Training language models to follow instructions
601 with human feedback. In *Advances in Neural Information*
602 *Processing Systems*, volume 35, 2022.
- 603 Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE*
604 *Transactions on Knowledge and Data Engineering*, 22,
2010.
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter,
S. Continual lifelong learning with neural networks: A
review. *Neural Networks*, 113, 2019.
- Paul, D. Asymptotics of sample eigenstructure for a large
dimensional spiked covariance model. *Statistica Sinica*,
2007.
- Plaut, E. From Principal Subspaces to Principal Com-
ponents with Linear Autoencoders. *arXiv preprint*
arXiv:1804.10253, 2018.
- Radhakrishnan, A., Belkin, M., and Uhler, C. Overparam-
eterized Neural Networks Implement Associative Memory.
Proceedings of the National Academy of Sciences, 117,
2020.
- Raposo, D., Kaufman, M. T., and Churchland, A. K. A
category-free neural population supports evolving de-
mands during decision-making. *Nature neuroscience*,
17, 2014.
- Refinetti, M. and Goldt, S. The dynamics of representation
learning in shallow, non-linear autoencoders. *Journal*
of Statistical Mechanics: Theory and Experiment, 2023,
2023.
- Richards, D., Mourada, J., and Rosasco, L. Asymptotics
of ridge(less) regression under general source condition.
In *International Conference on Artificial Intelligence and*
Statistics, 2020.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw,
N. D., Miller, E. K., and Fusi, S. The importance of mixed
selectivity in complex cognitive tasks. *Nature*, 497, 2013.
- Rubio, F. and Mestre, X. Spectral convergence for a general
class of random matrices. *Statistics & Probability Letters*,
81, 2011.
- Shima, K. and Tanji, J. Neuronal activity in the supple-
mentary and presupplementary motor areas for temporal
organization of multiple movements. *Journal of neuro-*
physiology, 84, 2000.

- 605 Song, Y., Bhattacharya, S., and Sur, P. Generalization error
606 of min-norm interpolators in transfer learning. *arXiv*
607 *preprint arXiv:2406.13944*, 2024.
- 608 Sonthalia, R., Murray, M., and Montúfar, G. F. Low rank
609 gradients and where to find them. In *Advances in Neural*
610 *Information Processing Systems*, 2025.
- 612 Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan,
613 D., and Duncan, J. Dynamic coding for cognitive control
614 in prefrontal cortex. *Neuron*, 78, 2013.
- 615 Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., and Hardt,
616 M. Test-time training with self-supervision for generaliza-
617 tion under distribution shifts. In *International conference*
618 *on machine learning*. PMLR, 2020.
- 620 Tahir, J., Ganguli, S., and Rotskoff, G. M. Features are
621 fate: a theory of transfer learning in high-dimensional
622 regression. In *Forty-second International Conference on*
623 *Machine Learning*, 2025.
- 624 Tang, E., Mattar, M. G., Giusti, C., Lydon-Staley, D. M.,
625 Thompson-Schill, S. L., and Bassett, D. S. Effective learn-
626 ing is accompanied by high-dimensional and efficient
627 representations of neural activity. *Nature neuroscience*,
628 22, 2019.
- 630 Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan
631 J. Tibshirani. Surprises in high-dimensional ridgeless
632 least squares interpolation. *The Annals of Statistics*, 50,
633 2022.
- 634 Tye, K., Miller, E., Taschbach, F., Benna, M., Rigotti, M.,
635 and Fusi, S. Mixed selectivity: Cellular computations for
636 complexity. *Neuron*, 112, 2024.
- 638 Valeriani, L., Doimo, D., Cuturello, F., Laio, A., Ansuini,
639 A., and Cazzaniga, A. The geometry of hidden representa-
640 tions of large transformer models. In *Advances in Neural*
641 *Information Processing Systems*, volume 36, 2023.
- 643 Vershynin, R. *High-Dimensional Probability: An Introduc-*
644 *tion with Applications in Data Science*. Cambridge Series
645 in Statistical and Probabilistic Mathematics. Cambridge
646 University Press, 2026.
- 647 Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Man-
648 zagol, P.-A. Stacked denoising autoencoders: Learning
649 useful representations in a deep network with a local de-
650 noising criterion. *Journal of Machine Learning Research*,
651 11, 2010.
- 653 Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and
654 Bowman, S. GLUE: A multi-task benchmark and analy-
655 sis platform for natural language understanding. In *Pro-*
656 *ceedings of the 2018 EMNLP Workshop BlackboxNLP:*
657 *Analyzing and Interpreting Neural Networks for NLP*,
658 2018.
- 659 Wang, B., Johnston, W. J., and Fusi, S. A mathemat-
ical theory for understanding when abstract represen-
tations emerge in neural networks. *arXiv preprint*
arXiv:2510.09816, 2026.
- Wu, D. and Xu, J. On the optimal weighted ℓ_2 regulariza-
tion in overparameterized linear regression. In *Advances*
in Neural Information Processing Systems, volume 33,
2020.
- Xu, J. and Hsu, D. J. On the number of variables to use in
principal component regression. In *Advances in Neural*
Information Processing Systems, volume 32, 2019.
- Yang, F., Zhang, H. R., Wu, S., Re, C., and Su, W. J. Precise
high-dimensional asymptotics for quantifying heteroge-
neous transfers. *Journal of Machine Learning Research*,
26, 2025.
- Yao, J., Bai, Z., and Zheng, S.-R. *Large sample covariance*
matrices and high-dimensional data analysis. Cambridge
series in statistical and probabilistic mathematics. Cam-
bridge university press, 2015.
- Zenke, F., Poole, B., and Ganguli, S. Continual learning
through synaptic intelligence. In *International Confer-*
ence on Machine Learning. PMLR, 2017.
- Zhang, T. and Hashimoto, T. B. On the inductive bias of
masked language modeling: From statistical to syntactic
dependencies. In *Proceedings of the 2021 Conference of*
the North American Chapter of the Association for Com-
putational Linguistics: Human Language Technologies,
2021.
- Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong,
H., and He, Q. A comprehensive survey on transfer
learning. *Proceedings of the IEEE*, PP, 2020.

Appendix

A. Theory

A.1. ASSUMPTIONS

Assumption .1 (High-dimensional asymptotic regime). Let the number of features be p , and n_l, n_u, m be such that

$$\frac{p}{n_l} \rightarrow \gamma_l \in (0, \infty), \quad \frac{p}{n_u} \rightarrow \gamma_u \in (0, \infty), \quad \frac{m}{p} \rightarrow \alpha \in [0, 1].$$

Furthermore, we define $\gamma_{\text{eff}} := \frac{m}{n_l} \rightarrow \alpha\gamma_l \in [0, \infty)$.

In words, α is the fraction of dimensions to keep, γ_u and γ_l are the aspect ratios the prior and downstream datasets, and γ_{eff} is the effective aspect ratio of the compressed downstream dataset.

Assumption .2 (Spiked covariance model). For each p , the population covariance matrix is

$$\Sigma = \mathbf{I}_p + (\lambda - 1)\mathbf{v}\mathbf{v}^\top,$$

where $\lambda > 1$ and the unit vector $\mathbf{v} \in \mathbb{R}^p$ are fixed.

Note that the eigendecomposition of Σ is

$$\Sigma = \mathbf{V}\text{diag}(\lambda, 1, \dots, 1)\mathbf{V}^\top,$$

with $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p] \in \mathbb{R}^{p \times p}$. As we often refer directly to the spike eigenvector, we keep the simpler notation $\mathbf{v}_1 = \mathbf{v}$.

Assumption .3 (Data generation). The input matrices $\mathbf{X}_l \in \mathbb{R}^{n_l \times p}$ and $\mathbf{X}_u \in \mathbb{R}^{n_u \times p}$ have i.i.d. rows distributed as

$$\mathbf{x} \sim \mathcal{N}(0, \Sigma).$$

Equivalently, there exists a matrix $\mathbf{Z} \in \mathbb{R}^{n_l \times p}$ with i.i.d. entries distributed as $\mathcal{N}(0, 1)$, such that $\mathbf{X}_l = \mathbf{Z}\Sigma^{1/2}$.

Assumption .4 (Task weight vector). Let the alignment between the true signal \mathbf{w}^* and the data covariance structure Σ be given by the spectral measure

$$\hat{G}_{\mathbf{w}^*}(\tau) = \frac{1}{\bar{w}^*} \sum_{i=1}^p (\mathbf{w}^{*\top} \mathbf{v}_i)^2 \mathbf{1}\{\lambda_i(\Sigma) \leq \tau\}.$$

Then, there exists a limiting spectrum $G_{\mathbf{w}^*}$ such that $\hat{G}_{\mathbf{w}^*} \Rightarrow G_{\mathbf{w}^*}$ almost surely.

Sample covariance eigenspectrum. We further note that sample covariance matrices $\mathbf{S}_u = \frac{1}{n_u} \mathbf{X}_u^\top \mathbf{X}_u$ and $\mathbf{S}_l = \frac{1}{n_l} \mathbf{X}_l^\top \mathbf{X}_l$ have empirical spectral measures

$$\hat{F}^{\mathbf{S}_u}(\theta) = \frac{1}{p} \sum_{i=1}^p \mathbf{1}\{\lambda_i(\mathbf{S}_u) \leq \theta\}, \quad \hat{F}^{\mathbf{S}_l}(\theta) = \frac{1}{p} \sum_{i=1}^p \mathbf{1}\{\lambda_i(\mathbf{S}_l) \leq \theta\},$$

which converge weakly almost surely to Marcenko-Pastur laws (Marčenko & Pastur, 1967; Bai & Silverstein, 2010) with aspect ratios γ_u and γ_l respectively:

$$\hat{F}^{\mathbf{S}_u} \Rightarrow F_{\gamma_u}, \quad \hat{F}^{\mathbf{S}_l} \Rightarrow F_{\gamma_l}.$$

For completeness, we state the Marcenko-Pastur density with an aspect ratio γ :

$$f_\gamma(\lambda) = \frac{\sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}}{2\pi\lambda\gamma} + \mathbf{1}_{\gamma > 1} \left(1 - \frac{1}{\gamma}\right) \delta(\lambda), \quad (7)$$

$$\lambda_+ = (1 + \sqrt{\gamma})^2,$$

$$\lambda_- = (1 - \sqrt{\gamma})^2,$$

with λ_+, λ_- denoting the edges of the support of the distribution.

A.2. MAIN IDEAS OF THE PROOFS

The key idea of each proof is to decompose the model's estimator $\hat{\mathbf{w}}$ into a signal term, a bias term arising from the component of \mathbf{w}^* outside the relevant subspace, and two variance terms arising from noise and model mismatch. The Gaussian structure of the data allows us to express the component of \mathbf{X}_l orthogonal to the prior subspace \mathbf{U}_m as a linear function of $\mathbf{X}_l \mathbf{U}_m$ plus independent noise \mathbf{E} . This reduces the analysis to understanding various projections onto the row space of spiked Gaussian matrices.

Conditional Gaussian decomposition under prior subspace projections. Let $\mathbf{X}_l \in \mathbb{R}^{n_l \times p}$ be a random matrix whose rows $\mathbf{x}_i \sim \mathcal{N}(0, \Sigma)$ are i.i.d., and let $\mathbf{U}_m \in \mathbb{R}^{p \times m}$ be a random matrix with orthonormal columns. Let $\mathbf{U}_\perp \in \mathbb{R}^{p \times (p-m)}$ be such that $[\mathbf{U}_m \ \mathbf{U}_\perp]$ is an orthogonal matrix. Define

$$\mathbf{X}_m = \mathbf{X}_l \mathbf{U}_m \in \mathbb{R}^{n_l \times m}, \quad \mathbf{X}_\perp = \mathbf{X}_l \mathbf{U}_\perp \in \mathbb{R}^{n_l \times (p-m)}.$$

Then, conditioning on \mathbf{U}_m , the rows of $(\mathbf{X}_m, \mathbf{X}_\perp)$ are jointly Gaussian with covariance

$$\begin{bmatrix} \Sigma_{mm} & \Sigma_{m\perp} \\ \Sigma_{\perp m} & \Sigma_{\perp\perp} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_m^\top \Sigma \mathbf{U}_m & \mathbf{U}_m^\top \Sigma \mathbf{U}_\perp \\ \mathbf{U}_\perp^\top \Sigma \mathbf{U}_m & \mathbf{U}_\perp^\top \Sigma \mathbf{U}_\perp \end{bmatrix}.$$

Moreover, due to the conditional distribution, \mathbf{X}_\perp given \mathbf{X}_m satisfies

$$\mathbf{X}_\perp = \mathbf{X}_m \Sigma_{mm}^{-1} \Sigma_{m\perp} + \mathbf{E}, \quad (8)$$

where \mathbf{E} is independent of \mathbf{X}_m and its rows are i.i.d. with distribution $\mathbf{e}_i \sim \mathcal{N}(0, \bar{\Sigma})$, with $\bar{\Sigma} := \Sigma_{\perp\perp} - \Sigma_{\perp m} \Sigma_{mm}^{-1} \Sigma_{m\perp}$. Under the spiked covariance model of Assumption .2, the population covariance Σ_{mm} of \mathbf{X}_m equals

$$\Sigma_{mm} = \mathbf{I}_m + (\tilde{\lambda} - 1) \tilde{\mathbf{v}} \tilde{\mathbf{v}}^\top, \quad (9)$$

where $\tilde{\lambda} = 1 + (\lambda - 1) \|\mathcal{P}_m \mathbf{v}\|^2$ and $\tilde{\mathbf{v}} = \frac{\mathbf{U}_m^\top \mathbf{v}}{\|\mathcal{P}_m \mathbf{v}\|}$. Furthermore, after some algebraic manipulations using the Sherman-Morrison formula, it follows that

$$\Sigma_{mm}^{-1} = \mathbf{I}_m - \frac{\tilde{\lambda} - 1}{\tilde{\lambda}} \tilde{\mathbf{v}} \tilde{\mathbf{v}}^\top, \quad (10)$$

$$\Sigma_{mm}^{-1} \Sigma_{m\perp} = \frac{(\lambda - 1) \|\mathcal{P}_m \mathbf{v}\|}{\tilde{\lambda}} \tilde{\mathbf{v}} \mathbf{v}^\top \mathbf{U}_\perp, \quad (11)$$

$$\bar{\Sigma} = \mathbf{I}_{p-m} + \frac{\lambda - 1}{\tilde{\lambda}} \mathbf{U}_\perp^\top \mathbf{v} \mathbf{v}^\top \mathbf{U}_\perp. \quad (12)$$

Decomposing the model's weight estimate. From the definition of $\hat{\mathbf{w}}$ in Equation (1) and the decomposition in Equation (8), we can write:

$$\begin{aligned} \hat{\mathbf{w}} &= (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathcal{P}_m \mathbf{w}^* + (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathcal{P}_\perp \mathbf{w}^* + (\mathbf{X}_l \mathcal{P}_m)^\dagger \xi \\ &= (\mathbf{X}_m \mathbf{U}_m^\top)^\dagger \mathbf{X}_m \mathbf{U}_m^\top \mathbf{w}^* + (\mathbf{X}_m \mathbf{U}_m^\top)^\dagger \mathbf{X}_m \Sigma_{mm}^{-1} \Sigma_{m\perp} \mathbf{U}_\perp^\top \mathbf{w}^* + (\mathbf{X}_m \mathbf{U}_m^\top)^\dagger \mathbf{E} \mathbf{U}_\perp^\top \mathbf{w}^* + (\mathbf{X}_m \mathbf{U}_m^\top)^\dagger \xi. \end{aligned}$$

Then, using Equation (11) leads to

$$\Sigma_{mm}^{-1} \Sigma_{m\perp} \mathbf{U}_\perp^\top \mathbf{w}^* = \underbrace{\frac{(\lambda - 1) (\mathbf{w}^{*\top} \mathcal{P}_\perp \mathbf{v})}{\tilde{\lambda}}}_{A} \mathbf{U}_m^\top \mathbf{v} \quad (13)$$

and by defining

$$\Pi_m := \mathbf{X}_m^\dagger \mathbf{X}_m \in \mathbb{R}^{m \times m} \quad \Pi := (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathcal{P}_m = \mathbf{U}_m \Pi_m \mathbf{U}_m^\top \in \mathbb{R}^{p \times p}$$

, we obtain

$$\begin{aligned} \hat{\mathbf{w}} &= \underbrace{\Pi \mathbf{w}^*}_{\text{signal from } \mathcal{P}_m \text{ subspace}} + \underbrace{A \cdot \Pi \mathbf{v}}_{\text{signal from } \mathcal{P}_\perp \text{ subspace that leaked into } \mathcal{P}_m} \\ &+ \underbrace{(\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{E} \mathbf{w}^*_\perp}_{\text{variance from model mismatch}} + \underbrace{(\mathbf{X}_l \mathcal{P}_m)^\dagger \xi}_{\text{variance from label noise}}, \end{aligned} \quad (14)$$

where $\mathbf{w}^*_{\perp} := \mathbf{U}_{\perp}^{\top} \mathbf{w}^*$. Note that the component of \mathbf{w}^* not involving label noise is

$$(\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{X}_l \mathbf{w}^* = \Pi \mathbf{w}^* + A \Pi \mathbf{v} + (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{E} \mathbf{w}^*_{\perp}. \quad (15)$$

The model's weight estimate decomposes into the four parts discussed below.

- $\Pi \mathbf{w}^*$: part of \mathbf{w}^* that lies in the subspace formed by \mathcal{P}_m and is identifiable from the available data. It is equal to $\mathcal{P}_m \mathbf{w}^*$ if $m < n_l$ (i.e., $\gamma_{\text{eff}} < 1$), since in that case $\Pi_m = \mathbf{I}_m$ (see Lemma .10); if $m > n_l$ (i.e., $\gamma_{\text{eff}} > 1$), this term causes bias due to the finite sample size.
- $A \Pi \mathbf{v}$: a part of \mathbf{w}^* that lies outside the subspace \mathcal{P}_m can leak into \mathcal{P}_m , causing the model to reconstruct it (in the \mathcal{P}_m subspace). The reason is the existence of the population covariance spike which makes the rows of $\mathbf{X}_l \mathcal{P}_{\perp}$ and $\mathbf{X}_l \mathcal{P}_m$ correlated, so the part outside the subspace leaks back into the subspace through this covariance structure.
- $(\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{E} \mathbf{w}^*_{\perp}$: the part of the \mathbf{w}^* that lies outside the subspace \mathcal{P}_m and does not leak into \mathcal{P}_m behaves like noise, eventually causing a double-descent-like curve that scales with $\|\mathbf{w}^*_{\perp}\|^2$.
- $(\mathbf{X}_l \mathcal{P}_m)^{\dagger} \xi$: label noise variance, causing another double-descent-like curve that scales with σ^2 .

A.3. PROOF FOR THE ESTIMATION ERROR

Proof of Theorem 4.1. The proof proceeds in three stages, exploiting the conditional independence structure of the data. First, we use the joint Gaussianity of \mathbf{X}_m and \mathbf{X}_{\perp} to rewrite \mathbf{X}_{\perp} in terms of \mathbf{X}_m and an independent noise matrix $\mathbf{E} \in \mathbb{R}^{n_l \times (p-m)}$ (see Section A.2). We then average over \mathbf{E} conditioned on \mathbf{X}_m and \mathbf{U}_m , then average over \mathbf{X}_m conditioned on \mathbf{U}_m , and finally over \mathbf{U}_m .

Step 1: Decomposition of the error. The estimation error splits into bias and variance:

$$\begin{aligned} \mathbb{E}_{\xi} [\|\mathbf{w}^* - \hat{\mathbf{w}}\|^2] &= \|(\mathbf{I} - (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{X}_l) \mathbf{w}^*\|^2 + \mathbb{E}_{\xi} [\xi^{\top} (\mathbf{X}_l \mathcal{P}_m)^{\dagger \top} (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \xi] \\ &\quad - 2 \mathbb{E}_{\xi} [\xi^{\top} (\mathbf{X}_l \mathcal{P}_m)^{\dagger \top} (\mathbf{I} - (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{X}_l) \mathbf{w}^*] \\ &= \|(\mathbf{I} - (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{X}_l) \mathbf{w}^*\|^2 + \mathbb{E}_{\xi} [\text{Tr} [\xi \xi^{\top} (\mathbf{X}_l \mathcal{P}_m)^{\dagger \top} (\mathbf{X}_l \mathcal{P}_m)^{\dagger}]] \\ &= \|(\mathbf{I} - (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{X}_l) \mathbf{w}^*\|^2 + \sigma^2 \text{Tr} [(\mathcal{P}_m \mathbf{X}_l^{\top} \mathbf{X}_l \mathcal{P}_m)^{\dagger}]. \end{aligned}$$

Substituting the result in (15) further decomposes the bias:

$$\begin{aligned} \frac{1}{\|\mathbf{w}^*\|^2} \|(\mathbf{I} - (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{X}_l) \mathbf{w}^*\|^2 &= \frac{1}{\|\mathbf{w}^*\|^2} \|\mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v} - (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{E} \mathbf{U}_{\perp}^{\top} \mathbf{w}^*\|^2 \\ &= \frac{1}{\|\mathbf{w}^*\|^2} \|\mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v}\|^2 + \frac{1}{\bar{w}^*} \|(\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{E} \mathbf{U}_{\perp}^{\top} \mathbf{w}^*\|^2 \\ &\quad - 2 \underbrace{\frac{1}{\|\mathbf{w}^*\|^2} (\mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v})^{\top} (\mathbf{X}_l \mathcal{P}_m)^{\dagger} \mathbf{E} \mathbf{U}_{\perp}^{\top} \mathbf{w}^*}_{\text{cross term } T}. \end{aligned}$$

Step 2: E-independent term. The first term $\|\mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v}\|^2$ does not depend on \mathbf{E} , so its conditional expectation on \mathbf{E} is itself. We expand using the property of the projection matrix Π that $\Pi^2 = \Pi$:

$$\frac{1}{\|\mathbf{w}^*\|^2} \|\mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v}\|^2 = 1 - \frac{\|\Pi \mathbf{w}^*\|^2}{\|\mathbf{w}^*\|^2} + \frac{A^2}{\|\mathbf{w}^*\|^2} \|\Pi \mathbf{v}\|^2, \quad (16)$$

where the cross term $-2A \mathbf{w}^{*\top} \Pi \mathbf{v} + 2A \mathbf{w}^{*\top} \Pi \Pi \mathbf{v} = 0$. Substituting $A = \frac{\lambda-1}{\lambda} (\mathbf{v}^{\top} \mathcal{P}_{\perp} \mathbf{w}^*)$ gives

$$1 - \frac{\|\Pi \mathbf{w}^*\|^2}{\|\mathbf{w}^*\|^2} + \frac{(\lambda-1)^2 (\mathbf{v}^{\top} \mathcal{P}_{\perp} \mathbf{w}^*)^2}{\tilde{\lambda}^2 \|\mathbf{w}^*\|^2} \|\Pi \mathbf{v}\|^2 \Rightarrow 1 - \bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*} + \frac{(\lambda-1)^2}{\lambda^2} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}^2 \bar{\pi}_{\mathbf{v}, \mathbf{v}} \text{ almost surely,} \quad (17)$$

825 which matches the first two terms of the bias. The projections $\frac{\|\Pi \mathbf{w}^*\|^2}{\|\mathbf{w}^*\|^2}$ and $\|\Pi \mathbf{v}\|^2$ converge almost surely to $\bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*}, \bar{\pi}_{\mathbf{v}, \mathbf{v}}$
 826 in Equation (41) according to Lemma .10. The quantity $\frac{\mathbf{v}^\top \mathcal{P}_\perp \mathbf{w}^*}{\|\mathbf{w}^*\|}$ converges almost surely to $\bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}$ in Equation (36)
 827 according to Lemma .9, and $\tilde{\lambda} = 1 + (\lambda - 1)\|\mathcal{P}_m \mathbf{v}\|^2$ converges to $\bar{\lambda}$ almost surely by the continuous mapping theorem and
 828 the almost sure convergence of $\|\mathcal{P}_m \mathbf{v}\|^2 \rightarrow \bar{p}_{\mathbf{v}, \mathbf{v}}$ from Equation (32) (see Lemma .8). By the continuous mapping theorem,
 829 the combined expression in (17) converges almost surely.
 830

831 **Step 3: Cross term linear in \mathbf{E} .** We write the cross term as
 832

$$T = \mathbf{c}^\top \mathbf{E} \mathbf{d},$$

833 with
 834

$$\begin{aligned} \mathbf{c} &:= \frac{1}{\|\mathbf{w}^*\|} (\mathbf{X}_l \mathcal{P}_m)^\dagger (\mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v}) \in \mathbb{R}^{n_l}, \\ \mathbf{d} &:= \frac{1}{\|\mathbf{w}^*\|} \mathbf{U}_\perp^\top \mathbf{w}^* \in \mathbb{R}^{p-m}. \end{aligned}$$

841 Conditioned on $(\mathbf{X}_m, \mathbf{U}_m)$, the vectors \mathbf{c} and \mathbf{d} are deterministic and independent of \mathbf{E} . Since the rows of \mathbf{E} are i.i.d.
 842 Gaussian with covariance $\bar{\Sigma}$, it follows that
 843

$$T \mid (\mathbf{X}_m, \mathbf{U}_m) \sim \mathcal{N}(0, (\mathbf{d}^\top \bar{\Sigma} \mathbf{d}) \|\mathbf{c}\|^2).$$

844 We now bound the variance. First,
 845

$$\mathbf{d}^\top \bar{\Sigma} \mathbf{d} = \frac{\|\mathcal{P}_\perp \mathbf{w}^*\|^2}{\|\mathbf{w}^*\|^2} + \frac{\lambda - 1}{\tilde{\lambda}} \frac{(\mathbf{w}^{*\top} \mathcal{P}_\perp \mathbf{v})^2}{\|\mathbf{w}^*\|^2},$$

846 which by the almost sure limits of $\|\mathcal{P}_\perp \mathbf{w}^*\|^2$, $\tilde{\lambda}$, and $\mathbf{w}^{*\top} \mathcal{P}_\perp \mathbf{v}$, converges almost surely to
 847

$$\bar{\sigma}_{\text{eff}}^2 = \bar{p}_{\perp, \mathbf{w}^*, \mathbf{w}^*} + \frac{\lambda - 1}{\bar{\lambda}} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}^2.$$

848 In particular, there exists a deterministic constant C_1 such that for all sufficiently large p ,
 849

$$\mathbf{d}^\top \bar{\Sigma} \mathbf{d} \leq C_1 \quad \text{almost surely.}$$

850 It therefore remains to control $\|\mathbf{c}\|^2$. We have
 851

$$\|\mathbf{c}\|^2 \leq \frac{1}{\|\mathbf{w}^*\|^2} \frac{\|\mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v}\|^2}{\sigma_{\min}(\mathbf{X}_m)^2}.$$

852 By Step 2, the numerator is bounded almost surely for all sufficiently large p . For the denominator, writing $\mathbf{X}_m = \mathbf{G} \Sigma_{mm}^{1/2}$
 853 (see Lemma .5) gives
 854

$$\sigma_{\min}(\mathbf{X}_m) \geq \sigma_{\min}(\mathbf{G}) \sigma_{\min}(\Sigma_{mm}^{1/2}).$$

855 Note that $\sigma_{\min}(\Sigma_{mm}^{1/2}) = \min\{1, \sqrt{\tilde{\lambda}}\}$ for any p . By the singular value inequality for Gaussian random matrices (see e.g.
 856 Exercise 7.13 in (Vershynin, 2026)), for all $t > 0$,
 857

$$\mathbb{P}(\sigma_{\min}(\mathbf{G}) \leq |\sqrt{n_l} - \sqrt{m}| - t) \leq 2e^{-t^2/2}.$$

858 Since $\gamma_{\text{eff}} = m/n_l \neq 1$, we have
 859

$$|\sqrt{n_l} - \sqrt{m}| = \sqrt{p} \frac{|1 - \sqrt{\gamma_{\text{eff}}}|}{\sqrt{\gamma_l}} = \sqrt{p} c_\gamma.$$

860 Choosing $t = \frac{1}{2} c_\gamma \sqrt{p}$ and setting $c = c_\gamma^2/8$, we obtain
 861

$$\mathbb{P}\left(\sigma_{\min}(\mathbf{G}) \leq \frac{1}{2} c_\gamma \sqrt{p}\right) \leq 2 \exp\left(-\frac{c_\gamma^2}{8} p\right).$$

Noting that $\tilde{\lambda} \leq \lambda$, and setting $a_\gamma := \frac{1}{4}c_\gamma^2 \min\{1, \lambda\}$, $b_\gamma := \frac{1}{8}c_\gamma^2$, we have

$$\mathbb{P}(\sigma_{\min}(\mathbf{X}_m)^2 \leq a_\gamma p) \leq 2e^{-b_\gamma p}.$$

Therefore, there exists $C_2 > 0$ such that

$$\mathbb{P}\left(\|\mathbf{c}\|^2 \leq \frac{C_2}{p}\right) \geq 1 - 2e^{-b_\gamma p}.$$

Define the event

$$\mathcal{E}_p := \left\{ \|\mathbf{c}\|^2 \leq \frac{C_2}{p} \right\}.$$

Then

$$\mathbb{P}(\mathcal{E}_p^c) \leq 2e^{-b_\gamma p}.$$

For any $\delta > 0$,

$$\mathbb{P}(|T| \geq \delta) = \mathbb{P}(|T| \geq \delta, \mathcal{E}_p^c) + \mathbb{P}(|T| \geq \delta, \mathcal{E}_p) \leq \mathbb{P}(\mathcal{E}_p^c) + \mathbb{P}(|T| \geq \delta, \mathcal{E}_p).$$

Noting that there exists C_1 s.t. $\mathbf{d}^\top \bar{\Sigma} \mathbf{d} \leq C_1$ for any large p , on \mathcal{E}_p ,

$$\text{Var}(T \mid \mathbf{X}_m, \mathbf{U}_m) \leq \frac{C_1 C_2}{p}.$$

Thus, on \mathcal{E}_p , we have

$$\mathbb{P}(|T| \geq \delta \mid \mathbf{X}_m, \mathbf{U}_m) \leq 2 \exp(-c_\delta p),$$

and since this holds for any $\mathbf{X}_m, \mathbf{U}_m$, by the law of total probability, we have

$$\mathbb{P}(|T| \geq \delta, \mathcal{E}_p) \leq 2e^{-c_\delta p}.$$

Therefore,

$$\mathbb{P}(|T| \geq \delta) \leq 2e^{-b_\gamma p} + 2e^{-c_\delta p} = Ce^{-cp}.$$

Since $\sum_p Ce^{-cp} < \infty$, the Borel–Cantelli lemma implies

$$T \xrightarrow{a.s.} 0.$$

Step 4: Average of the trace terms. The variance term contains

$$\text{Tr}[(\mathcal{P}_m \mathbf{X}_l^\top \mathbf{X}_l \mathcal{P}_m)^\dagger] = \text{Tr}[(\mathbf{X}_m^\top \mathbf{X}_m)^\dagger].$$

We note that the term $\text{Tr}[(\mathbf{X}_m \mathbf{X}_m^\top)^\dagger]$ is equivalent to the above, since the non-zero eigenvalues of $\mathbf{X}_m^\top \mathbf{X}_m$ and $\mathbf{X}_m \mathbf{X}_m^\top$ are the same. By Lemma .6, the trace term has a limiting value $\frac{\min\{1, \gamma_{\text{eff}}\}}{|\gamma_{\text{eff}} - 1|}$ almost surely.

Step 5: Quadratic term in \mathbf{E} . Define

$$Q := \mathbf{w}^{*\top} \mathbf{U}_\perp \mathbf{E}^\top (\mathbf{X}_m \mathbf{X}_m^\top)^\dagger \mathbf{E} \mathbf{U}_\perp \mathbf{w}^*.$$

Let

$$\mathbf{a} := \mathbf{U}_\perp^\top \mathbf{w}^*, \quad \mathbf{A} := (\mathbf{X}_m \mathbf{X}_m^\top)^\dagger.$$

Conditioned on $(\mathbf{X}_m, \mathbf{U}_m)$, the vector \mathbf{a} and the matrix \mathbf{A} are fixed and independent of \mathbf{E} . Since the rows of \mathbf{E} are i.i.d. Gaussian with covariance $\bar{\Sigma}$, the vector

$$\mathbf{c} := \mathbf{E} \mathbf{a}$$

is Gaussian with

$$\mathbf{c} \mid (\mathbf{X}_m, \mathbf{U}_m) \sim \mathcal{N}(0, (\mathbf{a}^\top \bar{\Sigma} \mathbf{a}) \mathbf{I}_{n_l}).$$

Therefore

$$Q = \mathbf{c}^\top \mathbf{A} \mathbf{c}.$$

First, the conditional expectation is

$$\mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m] = (\mathbf{a}^\top \bar{\Sigma} \mathbf{a}) \text{Tr}(\mathbf{A}).$$

Here

$$\mathbf{a}^\top \bar{\Sigma} \mathbf{a} = \mathbf{w}^{*\top} \mathbf{U}_\perp \bar{\Sigma} \mathbf{U}_\perp^\top \mathbf{w}^* = \|\mathbf{w}^*\|^2 \sigma_{\text{eff}}^2.$$

By the deterministic equivalents used above,

$$\|\mathbf{w}^*\|^2 \sigma_{\text{eff}}^2 \xrightarrow{a.s.} \bar{w}^* \bar{\sigma}_{\text{eff}}^2.$$

Moreover, by the pseudo-inverse trace computation in Step 4,

$$\text{Tr}[(\mathbf{X}_m \mathbf{X}_m^\top)^\dagger] \xrightarrow{a.s.} \frac{\min\{1, \gamma_{\text{eff}}\}}{|\gamma_{\text{eff}} - 1|}.$$

Thus

$$\mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m] \xrightarrow{a.s.} \bar{w}^* \bar{\sigma}_{\text{eff}}^2 \frac{\min\{1, \gamma_{\text{eff}}\}}{|\gamma_{\text{eff}} - 1|}.$$

It remains to show concentration around this conditional expectation. Write

$$\mathbf{c} = \sqrt{\mathbf{a}^\top \bar{\Sigma} \mathbf{a}} \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_{n_l}),$$

conditionally on $(\mathbf{X}_m, \mathbf{U}_m)$. Then

$$Q - \mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m] = (\mathbf{a}^\top \bar{\Sigma} \mathbf{a}) (\mathbf{z}^\top \mathbf{A} \mathbf{z} - \text{Tr}(\mathbf{A})).$$

By the Hanson–Wright inequality, for every fixed $\delta > 0$,

$$\mathbb{P}_{\mathbf{E}}(|Q - \mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m]| \geq \delta \mid \mathbf{X}_m, \mathbf{U}_m) \leq 2 \exp \left[-c \min \left(\frac{\delta^2}{(\mathbf{a}^\top \bar{\Sigma} \mathbf{a})^2 \|\mathbf{A}\|_F^2}, \frac{\delta}{(\mathbf{a}^\top \bar{\Sigma} \mathbf{a}) \|\mathbf{A}\|} \right) \right].$$

We now bound the quantities appearing in this inequality. As above,

$$\mathbf{a}^\top \bar{\Sigma} \mathbf{a} = \|\mathbf{w}^*\|^2 \sigma_{\text{eff}}^2$$

is bounded for all sufficiently large p almost surely. Furthermore, since $\gamma_{\text{eff}} \neq 1$, the nonzero spectrum of $\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m$ is bounded away from zero almost surely. Equivalently, the nonzero singular values of \mathbf{X}_m are of order \sqrt{p} (see argument in Step 3). Hence

$$\|\mathbf{A}\| = \|(\mathbf{X}_m \mathbf{X}_m^\top)^\dagger\| = O\left(\frac{1}{p}\right),$$

and

$$\|\mathbf{A}\|_F^2 = \sum_{i: \sigma_i(\mathbf{X}_m) > 0} \frac{1}{\sigma_i(\mathbf{X}_m)^4} = O\left(\frac{1}{p}\right).$$

Consequently, for all sufficiently large p ,

$$\mathbb{P}_{\mathbf{E}}(|Q - \mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m]| \geq \delta \mid \mathbf{X}_m, \mathbf{U}_m) \leq 2e^{-c\delta p}.$$

The right-hand side is summable in p , so Borel–Cantelli gives

$$Q - \mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m] \xrightarrow{a.s.} 0.$$

Combining this concentration statement with the limit of the conditional expectation yields

$$Q \xrightarrow{a.s.} \bar{w}^* \bar{\sigma}_{\text{eff}}^2 \frac{\min\{1, \gamma_{\text{eff}}\}}{|\gamma_{\text{eff}} - 1|}.$$

Combining. Summing the contributions from Steps 1–5 and applying the continuity mapping theorem concludes the proof. \square

A.4. PROOF FOR THE GENERALISATION ERROR

Proof of Theorem 4.2. We decompose the error and take averages over \mathbf{E} , \mathbf{X}_m , \mathbf{U}_m .

Step 1: Decomposition of the generalisation error. We start from the definition of the generalisation error

$$E^{\text{gen}} = \mathbb{E}_{\mathbf{x}, y, \xi} [\|y - \mathbf{x}^\top \hat{\mathbf{w}}\|^2 \mid \mathbf{w}^*, \mathbf{X}_l, \mathbf{X}_u]. \quad (18)$$

Using the data model $y = \mathbf{x}^\top \mathbf{w}^* + \xi$ and independence of the test noise, this becomes

$$\begin{aligned} E^{\text{gen}} &= \mathbb{E}_{\mathbf{x}, \xi} [(\mathbf{w}^* - \hat{\mathbf{w}})^\top \mathbf{x} \mathbf{x}^\top (\mathbf{w}^* - \hat{\mathbf{w}}) \mid \mathbf{w}^*, \mathbf{X}_l, \mathbf{X}_u] + \sigma^2 \\ &= \mathbb{E}_\xi [(\mathbf{w}^* - \hat{\mathbf{w}})^\top \Sigma (\mathbf{w}^* - \hat{\mathbf{w}})] + \sigma^2. \end{aligned} \quad (19)$$

Using the spiked covariance model $\Sigma_{\text{new}} = \mathbf{I}_p + \sum_i (\nu_i - 1) \mathbf{v}_{\text{new}, i} \mathbf{v}_{\text{new}, i}^\top$, we obtain

$$E^{\text{gen}} = E^{\text{est}} + \sum_i (\nu_i - 1) \mathbb{E}_\xi [((\mathbf{w}^* - \hat{\mathbf{w}})^\top \mathbf{v}_{\text{new}, i})^2] + \sigma^2. \quad (20)$$

Substituting $\hat{\mathbf{w}} = (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathbf{w}^* + (\mathbf{X}_l \mathcal{P}_m)^\dagger \xi$ and expanding the quadratic form, the cross term vanishes in ξ expectation and we obtain a bias contribution

$$((\mathbf{w}^* - (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathbf{w}^*)^\top \mathbf{v}_{\text{new}, i})^2 \quad (21)$$

and a variance contribution proportional to

$$\mathbf{v}_{\text{new}, i}^\top (\mathcal{P}_m \mathbf{X}_l^\top \mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new}, i}. \quad (22)$$

Using the same decomposition and conditional Gaussian argument as in Equation (15), $(\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathbf{w}^* = \Pi \mathbf{w}^* + A \Pi \mathbf{v} + (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{E} \mathbf{w}^*_{\perp}$, we further decompose the bias

$$(\mathbf{w}^{*\top} - (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{X}_l \mathbf{w}^*)^\top \mathbf{v}_{\text{new}, i} = \mathbf{w}^{*\top} \mathbf{v}_{\text{new}, i} - \mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{w}^* - A \mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{v} - \mathbf{v}_{\text{new}, i}^\top (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{E} \mathbf{w}^*_{\perp}, \quad (23)$$

with the squared term being:

$$\begin{aligned} &\underbrace{(\mathbf{w}^{*\top} \mathbf{v}_{\text{new}, i} - \mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{w}^* - \frac{\lambda - 1}{\tilde{\lambda}} (\mathbf{w}^{*\top} \mathcal{P}_{\perp} \mathbf{v}) (\mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{v}))^2}_{\mathbf{E}\text{-independent}} + \underbrace{(\mathbf{v}_{\text{new}, i}^\top (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{E} \mathbf{w}^*_{\perp})^2}_{\text{quadratic in } \mathbf{E}} \\ &- 2 \underbrace{(\mathbf{w}^{*\top} \mathbf{v}_{\text{new}, i} - \mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{w}^* - \frac{\lambda - 1}{\tilde{\lambda}} (\mathbf{w}^{*\top} \mathcal{P}_{\perp} \mathbf{v}) (\mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{v})) (\mathbf{v}_{\text{new}, i}^\top (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{E} \mathbf{w}^*_{\perp})}_{\text{cross term}}. \end{aligned}$$

Step 2: E-independent term. The term linear in \mathbf{E} is made up of terms that have almost sure limits. In particular:

$$\begin{aligned} \mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{w}^* &\xrightarrow{a.s.} \sqrt{\bar{w}^*} \bar{\pi}_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} && \text{by Lemma .11} \\ \mathbf{v}_{\text{new}, i}^\top \Pi \mathbf{v} &\xrightarrow{a.s.} \bar{\pi}_{\mathbf{v}, \mathbf{v}_{\text{new}, i}} && \text{by Lemma .11} \\ \mathbf{w}^{*\top} \mathcal{P}_{\perp} \mathbf{v} &\xrightarrow{a.s.} \sqrt{\bar{w}^*} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}} && \text{by Lemma .9} \\ \tilde{\lambda} &\xrightarrow{a.s.} \bar{\lambda} && \text{by Lemma .8} \\ \mathbf{w}^{*\top} \mathbf{v}_{\text{new}, i} &\xrightarrow{a.s.} \sqrt{\bar{w}^*} \rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} && \text{by model assumptions} \end{aligned}$$

Thus, by the continuous mapping theorem, we have the almost sure convergence of the entire \mathbf{E} -independent term to

$$\bar{w}^* \left(\rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \bar{\pi}_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \frac{\lambda - 1}{\bar{\lambda}} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}} \bar{\pi}_{\mathbf{v}, \mathbf{v}_{\text{new}, i}} \right)^2.$$

1045 **Step 3: Cross term linear in \mathbf{E} .** The linear \mathbf{E} term has the form

$$1046 \quad T = \mathbf{c}^\top \mathbf{E} \mathbf{d},$$

1047 where

$$1048 \quad \mathbf{c} = C (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new},i}, \quad \mathbf{d} = \mathbf{w}^*_{\perp},$$

1049 and C is the deterministic prefactor defined in Step 2, which converges almost surely to an $O(1)$ limit.

1050 This term has the same structure as the term linear in \mathbf{E} in Theorem 4.1, Step 3, with the vector

$$1051 \quad \mathbf{w}^* - \Pi \mathbf{w}^* - A \Pi \mathbf{v}$$

1052 replaced by the unit vector $\mathbf{v}_{\text{new},i}$ and multiplied by the scalar $C = O(1)$. The same arguments apply and we conclude that

$$1053 \quad T \xrightarrow{a.s.} 0.$$

1054 **Step 4: Quadratic term in \mathbf{E} .** Define

$$1055 \quad Q := \mathbf{w}^*_{\perp}{}^\top \mathbf{E}^\top (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new},i} \mathbf{v}_{\text{new},i}^\top (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{E} \mathbf{w}^*_{\perp}.$$

1056 Let

$$1057 \quad \mathbf{a} := \mathbf{w}^*_{\perp}, \quad \mathbf{A} := (\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new},i} \mathbf{v}_{\text{new},i}^\top (\mathbf{X}_l \mathcal{P}_m)^\dagger.$$

1058 Conditioned on $(\mathbf{X}_m, \mathbf{U}_m)$, the vector \mathbf{a} and matrix \mathbf{A} are independent of \mathbf{E} , so $Q = \mathbf{a}^\top \mathbf{E}^\top \mathbf{A} \mathbf{E} \mathbf{a}$ is a quadratic Gaussian form.

1059 This is exactly the same object as in Theorem 4.1, Step 5, with a different choice of \mathbf{A} . The concentration argument therefore carries over verbatim once we verify that the same norm bounds hold.

1060 First, as in Theorem 4.1, Step 3,

$$1061 \quad \mathbf{a}^\top \bar{\Sigma} \mathbf{a} = O(1) \quad \text{almost surely.}$$

1062 Next, since \mathbf{A} is rank one,

$$1063 \quad \|\mathbf{A}\| = \|(\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new},i}\|^2, \quad \|\mathbf{A}\|_F^2 = \|\mathbf{A}\|^2.$$

1064 Using the bounds on $\|(\mathbf{X}_l \mathcal{P}_m)^\dagger\|_{\text{op}}$ established in Theorem 4.1, Step 3, together with $\|\mathbf{v}_{\text{new},i}\| = 1$, we obtain

$$1065 \quad \|\mathbf{A}\| = O\left(\frac{1}{p}\right), \quad \|\mathbf{A}\|_F^2 = O\left(\frac{1}{p^2}\right), \quad \text{almost surely.}$$

1066 The conditions required for the Hanson–Wright concentration used in Theorem 4.1, Step 5 are satisfied. Therefore, by the same argument as in Theorem 4.1, Step 5,

$$1067 \quad Q - \mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m] \xrightarrow{a.s.} 0.$$

1068 Finally, the conditional expectation is

$$1069 \quad \mathbb{E}_{\mathbf{E}}[Q \mid \mathbf{X}_m, \mathbf{U}_m] = (\mathbf{a}^\top \bar{\Sigma} \mathbf{a}) \text{Tr}(\mathbf{A}) = O\left(\frac{1}{p}\right) \xrightarrow{a.s.} 0$$

1070 **Step 5: Variance contribution.** Finally, we consider the term from Equation (22) and note it is the same as the norm of the matrix \mathbf{A} from Step 4:

$$1071 \quad \mathbf{v}_{\text{new},i}^\top (\mathcal{P}_m \mathbf{X}_l^\top \mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new},i} = \|(\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new},i}\|^2 = \|\mathbf{A}\| = O\left(\frac{1}{p}\right).$$

1072 Thus, this term converges to 0:

$$1073 \quad \mathbf{v}_{\text{new},i}^\top (\mathcal{P}_m \mathbf{X}_l^\top \mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{v}_{\text{new},i} \xrightarrow{a.s.} 0.$$

Combining. Thus, in the high-dimensional limit the only additional contribution relative to the estimation error is the spike bias term, and the generalisation error converges to

$$E_\infty^{\text{gen}} = E_\infty^{\text{est}} + \bar{w}^* \sum_i (\nu_i - 1) \left(\rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \bar{\pi}_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \frac{\lambda - 1}{\lambda} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}} \bar{\pi}_{\mathbf{v}, \mathbf{v}_{\text{new}, i}} \right)^2 + \sigma^2.$$

□

A.5. PROOF FOR THE TRAINING ERROR

Proof of Theorem 4.3. Let

$$\mathbf{H} = (\mathbf{X}_l \mathcal{P}_m)(\mathbf{X}_l \mathcal{P}_m)^\dagger$$

and recall that the target vector is $\mathbf{y}_d = \mathbf{X}_l \mathbf{w}^* + \xi$. Then, the residual vector is:

$$\begin{aligned} \mathbf{y}_d - \mathbf{X}_l \hat{\mathbf{w}} &= \mathbf{y}_d - (\mathbf{X}_l \mathcal{P}_m)(\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{y}_d - (\mathbf{X}_l \mathcal{P}_\perp)(\mathbf{X}_l \mathcal{P}_m)^\dagger \mathbf{y}_d \\ &= (\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{y}_d \\ &= (\mathbf{I}_{n_l} - \mathbf{H})(\mathbf{X}_l \mathbf{w}^* + \xi), \end{aligned} \tag{24}$$

where we note that the term multiplying two projections is 0:

$$(\mathbf{X}_l \mathcal{P}_\perp)(\mathbf{X}_l \mathcal{P}_m)^\dagger = \mathbf{X}_l \mathbf{U}_\perp \underbrace{\mathbf{U}_\perp^\top \mathbf{U}_m}_0 (\mathbf{X}_l \mathbf{U}_m)^\dagger = 0.$$

The training error splits into two terms since the noise ξ is independent and zero-mean:

$$\begin{aligned} E^{\text{train}} &= \frac{1}{n_l} \mathbb{E}_\xi \left[\|(\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{X}_l \mathbf{w}^* + (\mathbf{I}_{n_l} - \mathbf{H}) \xi\|^2 \right] \\ &= \underbrace{\frac{1}{n_l} \|(\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{X}_l \mathbf{w}^*\|^2}_{\text{Bias}^2} + \underbrace{\frac{1}{n_l} \mathbb{E}_\xi \left[\|(\mathbf{I}_{n_l} - \mathbf{H}) \xi\|^2 \right]}_{\text{Variance}}. \end{aligned} \tag{25}$$

We start by computing the bias. Using the decomposition $\mathbf{X}_l = \mathbf{X}_l \mathcal{P}_m + \mathbf{X}_l \mathcal{P}_\perp$ and the fact that $(\mathbf{I} - \mathbf{H}) \mathbf{X}_l \mathcal{P}_m = 0$, the bias simplifies to:

$$\text{Bias}^2 = \frac{1}{n_l} \|(\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{X}_l \mathcal{P}_\perp \mathbf{w}^*\|^2. \tag{26}$$

Substituting the conditional Gaussian decomposition, we have

$$\mathbf{X}_l \mathcal{P}_\perp = \mathbf{X}_m \boldsymbol{\Sigma}_{mm}^{-1} \boldsymbol{\Sigma}_{m\perp} \mathbf{U}_\perp^\top + \mathbf{E} \mathbf{U}_\perp^\top,$$

where \mathbf{E} is independent of \mathbf{X}_m . Thus, we obtain

$$\begin{aligned} (\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{X}_l \mathcal{P}_\perp \mathbf{w}^* &= (\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{X}_m \boldsymbol{\Sigma}_{mm}^{-1} \boldsymbol{\Sigma}_{m\perp} \mathbf{U}_\perp^\top \mathbf{w}^* + (\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{E} \mathbf{U}_\perp^\top \mathbf{w}^* \\ &= (\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{E} \mathbf{w}^*_\perp, \end{aligned} \tag{27}$$

where we used $(\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{X}_m = 0$ and defined $\mathbf{w}^*_\perp = \mathbf{U}_\perp^\top \mathbf{w}^*$. Note that the entries of $\mathbf{E} \mathbf{w}^*_\perp$ are i.i.d. $\mathcal{N}(0, \|\mathbf{w}^*\|^2 \sigma_{\text{eff}}^2)$, with

$$\sigma_{\text{eff}}^2 = \frac{\|\mathcal{P}_\perp \mathbf{w}^*\|^2}{\|\mathbf{w}^*\|^2} + \frac{\lambda - 1}{\lambda} \frac{(\mathbf{v}^\top \mathcal{P}_\perp \mathbf{w}^*)^2}{\|\mathbf{w}^*\|^2}. \tag{28}$$

Conditioned on \mathbf{X}_m , \mathbf{H} is deterministic, and the term $\|(\mathbf{I}_{n_l} - \mathbf{H}) \mathbf{E} \mathbf{w}^*_\perp\|^2$ is a quadratic form $z^\top (\mathbf{I}_{n_l} - \mathbf{H}) z$ where $z \sim \mathcal{N}(0, \|\mathbf{w}^*\|^2 \sigma_{\text{eff}}^2 \mathbf{I}_{n_l})$. Hence, we have

$$\mathbb{E}_{\mathbf{E}} [\text{Bias}^2 \mid \mathbf{X}_m] = \|\mathbf{w}^*\|^2 \frac{\sigma_{\text{eff}}^2}{n_l} \text{Tr}(\mathbf{I}_{n_l} - \mathbf{H}) = \|\mathbf{w}^*\|^2 \sigma_{\text{eff}}^2 \left(1 - \frac{\min(m, n_l)}{n_l} \right), \tag{29}$$

where the last equality holds almost surely (w.r.t. the randomness in \mathbf{X}_m).

Next, to compute the variance, recall that the label noise ξ is isotropic $\mathcal{N}(0, \sigma^2 \mathbf{I}_{n_l})$. Thus, we have

$$\frac{1}{n_l} \mathbb{E}_\xi [\xi^\top (\mathbf{I}_{n_l} - \mathbf{H}) \xi] = \frac{\sigma^2}{n_l} \text{Tr}(\mathbf{I}_{n_l} - \mathbf{H}) = \sigma^2 \left(1 - \frac{\min(m, n_l)}{n_l} \right), \quad (30)$$

where the last equality again holds almost surely (w.r.t. the randomness in \mathbf{X}_m). Note that

$$1 - \frac{\min(m, n_l)}{n_l} = \max(0, 1 - \gamma_{\text{eff}}).$$

To conclude, it suffices to recall from the proof of Theorem 4.1 that $\sigma_{\text{eff}}^2 \xrightarrow{\text{a.s.}} \bar{\sigma}_{\text{eff}}^2$, and the argument is complete. \square

A.6. AUXILIARY LEMMAS

Lemma .5 (Limiting spectral distribution of \mathbf{X}_m). *Consider the setup of Section 3, and define*

$$\mathbf{X}_m := \mathbf{X}_l \mathbf{U}_m.$$

Then the empirical spectral measure of $\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m$ converges weakly almost surely to the Marchenko–Pastur law with aspect ratio γ_{eff} , i.e.,

$$\hat{F}_{\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m} \Rightarrow F_{\gamma_{\text{eff}}} \quad \text{almost surely.}$$

Proof. Fix a realization of the full sequence $(\mathbf{U}_m)_{p \geq 1}$. Conditional on this realization, the rows of $\mathbf{X}_m = \mathbf{X}_l \mathbf{U}_m$ are i.i.d. Gaussian in \mathbb{R}^m with covariance

$$\Sigma_{mm} := \mathbf{U}_m^\top \Sigma \mathbf{U}_m.$$

Using the form of Σ from Assumption .2, we obtain

$$\Sigma_{mm} = \mathbf{U}_m^\top \left(\mathbf{I}_p + (\lambda - 1) \mathbf{v} \mathbf{v}^\top \right) \mathbf{U}_m = \mathbf{I}_m + (\lambda - 1) (\mathbf{U}_m^\top \mathbf{v}) (\mathbf{U}_m^\top \mathbf{v})^\top.$$

Hence, $\Sigma_{mm} - \mathbf{I}_m$ has rank at most one. Since $\Sigma_{mm} \succ 0$, define $\mathbf{G} \in \mathbb{R}^{n_l \times m}$ and $\mathbf{W} \in \mathbb{R}^{m \times m}$ by

$$\mathbf{G} := \mathbf{X}_m \Sigma_{mm}^{-1/2} \quad \mathbf{W} := \frac{1}{n_l} \mathbf{G}^\top \mathbf{G}.$$

Conditional on \mathbf{U}_m , \mathbf{G} has i.i.d. $\mathcal{N}(0, 1)$ entries. The sample covariance can be rewritten:

$$\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m = \Sigma_{mm}^{1/2} \left(\frac{1}{n_l} \mathbf{G}^\top \mathbf{G} \right) \Sigma_{mm}^{1/2} = \Sigma_{mm}^{1/2} \mathbf{W} \Sigma_{mm}^{1/2}.$$

Since $\Sigma_{mm} = \mathbf{I}_m + \mathbf{R}$ with $\text{rank}(\mathbf{R}) \leq 1$, the spectral decomposition shows that $\Sigma_{mm}^{1/2} - \mathbf{I}_m$ also has rank at most one. Consequently,

$$\Sigma_{mm}^{1/2} \mathbf{W} \Sigma_{mm}^{1/2} - \mathbf{W} = (\Sigma_{mm}^{1/2} - \mathbf{I}_m) \mathbf{W} \Sigma_{mm}^{1/2} + \mathbf{W} (\Sigma_{mm}^{1/2} - \mathbf{I}_m),$$

and thus

$$\text{rank} \left(\Sigma_{mm}^{1/2} \mathbf{W} \Sigma_{mm}^{1/2} - \mathbf{W} \right) \leq 2.$$

Let $\hat{F}^{\mathbf{A}}$ denote the empirical spectral distribution function of a real symmetric matrix \mathbf{A} . By the rank inequality (e.g., see Theorem A.43 in (Bai & Silverstein, 2010)),

$$\sup_{\theta \in \mathbb{R}} \left| \hat{F}_{\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m}(\theta) - \hat{F}^{\mathbf{W}}(\theta) \right| \leq \frac{2}{m} \xrightarrow{p \rightarrow \infty} 0.$$

Since \mathbf{W} is a standard Wishart matrix of size $m \times m$ with aspect ratio $\frac{m}{n_l} \rightarrow \gamma_{\text{eff}}$, by the Marchenko–Pastur theorem (Marčenko & Pastur, 1967),

$$\hat{F}^{\mathbf{W}} \Rightarrow F_{\gamma_{\text{eff}}} \quad \text{almost surely, conditional on } (\mathbf{U}_m)_{p \geq 1}.$$

Let $A := \left\{ F^{\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m} \Rightarrow F_{\gamma_{\text{eff}}} \right\}$. Taken together, the arguments above shows that:

$$\mathbb{P}(A \mid (\mathbf{U}_m)_{p \geq 1}) = 1.$$

We now remove the conditioning on $(\mathbf{U}_m)_{p \geq 1}$. By the law of total probability,

$$\mathbb{P}(A) = \mathbb{E}[\mathbb{P}(A \mid (\mathbf{U}_m)_{p \geq 1})] = \mathbb{E}[1] = 1.$$

This proves that

$$\hat{F}^{\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m} \Rightarrow F_{\gamma_{\text{eff}}} \quad \text{almost surely.}$$

□

Lemma .6 (Deterministic equivalent for the pseudoinverse trace). *Consider the setup of Section 3, and define*

$$\mathbf{X}_m := \mathbf{X}_l \mathbf{U}_m.$$

Then

$$\text{Tr}[(\mathbf{X}_m^\top \mathbf{X}_m)^\dagger] \xrightarrow{a.s.} \begin{cases} \frac{\gamma_{\text{eff}}}{1 - \gamma_{\text{eff}}}, & \gamma_{\text{eff}} < 1, \\ \frac{1}{\gamma_{\text{eff}} - 1}, & \gamma_{\text{eff}} > 1. \end{cases}$$

Proof. We have

$$\text{Tr}[(\mathbf{X}_m^\top \mathbf{X}_m)^\dagger] = \frac{1}{n_l} \text{Tr}\left(\left(\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m\right)^\dagger\right) = \frac{m}{n_l} \cdot \frac{1}{m} \text{Tr}\left(\left(\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m\right)^\dagger\right). \quad (31)$$

By Lemma .5, the empirical spectral distribution of $\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m$ converges weakly almost surely to the Marcenko–Pastur law with aspect ratio γ_{eff} . If $\gamma_{\text{eff}} < 1$, its support is bounded away from zero, and the claim follows by continuity of $x \mapsto 1/x$. If $\gamma_{\text{eff}} > 1$, there is an additional atom at zero; however, the pseudo-inverse removes this null space contribution, and the nonzero spectrum remains bounded away from zero, so the same argument applies on the effective support. Therefore,

$$\frac{1}{m} \text{Tr}\left(\left(\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m\right)^\dagger\right) = \int \frac{1}{x} dF^{\frac{1}{n_l} \mathbf{X}_m^\top \mathbf{X}_m}(x) \xrightarrow{a.s.} \int \frac{1}{x} dF_{\gamma_{\text{eff}}}(x).$$

Recall that

$$\int \frac{1}{x} dF_\gamma = \begin{cases} \frac{1}{1 - \gamma}, & \gamma < 1, \\ \frac{1}{\gamma(\gamma - 1)}, & \gamma > 1. \end{cases}$$

Substituting into (31) and using $\gamma = \gamma_{\text{eff}}$ gives

$$\text{Tr}[(\mathbf{X}_m^\top \mathbf{X}_m)^\dagger] \xrightarrow{a.s.} \begin{cases} \frac{\gamma_{\text{eff}}}{1 - \gamma_{\text{eff}}}, & \gamma_{\text{eff}} < 1, \\ \frac{1}{\gamma_{\text{eff}} - 1}, & \gamma_{\text{eff}} > 1, \end{cases}$$

which concludes the proof. □

Lemma .7 (Spectral measure of the overlap). *Consider the setup of Section 3, and let $(\mathbf{u}_i)_{i=1}^p$ be the eigenvectors of $\mathbf{S}_u = \frac{1}{n_u} \mathbf{X}_u^\top \mathbf{X}_u$. Let $\mathbf{a} \in \mathbb{R}^p$ be a fixed vector with population mass distribution $G_{\mathbf{a}}$ on the eigenvectors of Σ , i.e., $G_{\mathbf{a}} = \frac{(\mathbf{a}^\top \mathbf{v})^2}{\|\mathbf{a}\|^2} \delta_\lambda + \left(1 - \frac{(\mathbf{a}^\top \mathbf{v})^2}{\|\mathbf{a}\|^2}\right) \delta_1$. Define the empirical CDF*

$$\hat{G}_{\mathbf{a}}^{\text{sample}}(\theta) = \frac{1}{\|\mathbf{a}\|^2} \sum_{i=1}^p (\mathbf{a}^\top \mathbf{u}_i)^2 \mathbf{1}\{\theta \leq \lambda_i(\mathbf{S}_u)\}.$$

Then, $\hat{G}_{\mathbf{a}}^{\text{sample}} \Rightarrow G_{\mathbf{a}}^{\text{sample}}$ almost surely, where $G_{\mathbf{a}}^{\text{sample}}$ is such that (i) its support satisfies $\text{supp}(G_{\mathbf{a}}^{\text{sample}}) \subseteq \mathcal{S}_{\gamma_u} \cup \{\theta^*\} \cup \{0\}$ with \mathcal{S}_{γ_u} being the support of the Marcenko-Pastur density with aspect ratio γ_u ; (ii) it has a continuous density $g_{\mathbf{a}}^{\text{sample}}(\theta)$ on the interior of \mathcal{S}_{γ_u}

$$g_{\mathbf{a}}^{\text{sample}}(\theta) = \left(\int \frac{\tau \gamma_u}{\tau^2 - \tau(1 - \gamma_u + \theta) + \theta} dG_{\mathbf{a}}(\tau) \right) f_{\gamma_u}(\theta);$$

(iii) it has an atom at $\{\theta^*\}$ where $\theta^* = \frac{\lambda(\lambda - 1 + \gamma_u)}{\lambda - 1}$ is the location of the spike λ in the sample covariance spectrum

$$G_{\mathbf{a}}^{\text{sample}}(\{\theta^*\}) = \frac{1 - \frac{\gamma_u}{(\lambda-1)^2}}{1 + \frac{\gamma_u}{(\lambda-1)}} G_{\mathbf{a}}(\{\lambda\});$$

and (iv) it may have an atom at 0 of weight

$$G_{\mathbf{a}}^{\text{sample}}(\{0\}) = \begin{cases} G_{\mathbf{a}}(\{0\}) & \gamma_u < 1 \\ \int \frac{\gamma_u - 1}{\gamma_u - 1 + \tau} dG_{\mathbf{a}}(\tau), & \gamma_u > 1. \end{cases}$$

Proof. The result follows by specializing Lemma 3 in (Green & Romanov, 2025) to the spiked covariance model of Assumption .2. \square

Lemma .8 (Deterministic limit of $\|\mathcal{P}_m \mathbf{a}\|^2, \|\mathcal{P}_{\perp} \mathbf{a}\|^2$). Consider the setup of Section 3, and let $\mathbf{a} \in \mathbb{R}^p$ be a fixed vector with population mass distribution $G_{\mathbf{a}}$ on the eigenvectors of Σ . Let \mathcal{P}_m be the projection onto the top m eigenvectors of the prior sample covariance \mathbf{S}_u , with $m/p \rightarrow \alpha$. Let $G_{\mathbf{a}}^{\text{sample}}$ denote the deterministic limiting weighted spectral measure defined in Lemma .7. Let λ_t be the deterministic bulk threshold corresponding to retaining an asymptotic fraction α of principal components

$$F_{\gamma_u}(\lambda_t) = 1 - \alpha.$$

Then, almost surely,

$$\frac{\|\mathcal{P}_{\perp} \mathbf{a}\|^2}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} \begin{cases} \int_{\lambda_-}^{\lambda_t} g_{\mathbf{a}}^{\text{sample}}(\lambda), & \gamma_u < 1, \\ G_{\mathbf{a}}^{\text{sample}}(\{0\}) + \int_{\lambda_-}^{\lambda_t} g_{\mathbf{a}}^{\text{sample}}(\lambda), & \gamma_u > 1, \alpha < \frac{1}{\gamma_u}, \\ \frac{1 - \alpha}{1 - \frac{1}{\gamma_u}} G_{\mathbf{a}}^{\text{sample}}(\{0\}), & \gamma_u > 1, \alpha \geq \frac{1}{\gamma_u}, \end{cases} \quad (32)$$

where $\lambda_- = (1 - \sqrt{\gamma_u})^2$ is the infimum of the support of the Marcenko-Pastur distribution as defined in Equation (7) and

$$\|\mathcal{P}_m \mathbf{a}\|^2 = \|\mathbf{a}\|^2 - \|\mathcal{P}_{\perp} \mathbf{a}\|^2. \quad (33)$$

Proof. Let

$$\hat{G}_{\mathbf{a}}^{\text{sample}}(\theta) = \frac{1}{\|\mathbf{a}\|^2} \sum_{i=1}^p (\mathbf{a}^{\top} \mathbf{u}_i)^2 \mathbf{1}\{\theta \leq \lambda_i(\mathbf{S}_u)\}$$

be the weighted empirical spectral measure of \mathbf{S}_u . By Lemma .7, $\hat{G}_{\mathbf{a}}^{\text{sample}}$ converges almost surely to the deterministic measure $G_{\mathbf{a}}^{\text{sample}}$.

Since $\|\mathcal{P}_{\perp} \mathbf{a}\|^2 / \|\mathbf{a}\|^2$ is exactly the mass of $\hat{G}_{\mathbf{a}}^{\text{sample}}$ on the discarded eigenspaces, its limit is obtained by integrating $G_{\mathbf{a}}^{\text{sample}}$ over the discarded part of the spectrum.

If $\gamma_u < 1$, there is no nullspace, so only bulk directions are discarded. This gives

$$\frac{\|\mathcal{P}_{\perp} \mathbf{a}\|^2}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} \int_{\lambda_-}^{\lambda_t} g_{\mathbf{a}}^{\text{sample}}(\lambda).$$

If $\gamma_u > 1$ and $\alpha < 1/\gamma_u$, fewer than all positive-eigenvalue directions are retained, so the entire nullspace is discarded as well. Hence

$$\frac{\|\mathcal{P}_\perp \mathbf{a}\|^2}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} G_{\mathbf{a}}^{\text{sample}}(\{0\}) + \int_{\lambda_-}^{\lambda_t} g_{\mathbf{a}}^{\text{sample}}(\lambda).$$

Finally, if $\gamma_u > 1$ and $\alpha \geq 1/\gamma_u$, all positive-eigenvalue directions are retained, and the retained subspace is completed by adding

$$m - \text{rank}(\mathbf{S}_u) = m - n_u$$

orthonormal directions chosen uniformly at random in the nullspace, independently of \mathbf{a} . Therefore only the fraction

$$\frac{p - m}{p - n_u} \xrightarrow{a.s.} \frac{1 - \alpha}{1 - \frac{1}{\gamma_u}}$$

of the nullspace remains discarded. Since the total asymptotic mass of \mathbf{a} in the nullspace is $G_{\mathbf{a}}^{\text{sample}}(\{0\})$, we obtain

$$\frac{\|\mathcal{P}_\perp \mathbf{a}\|^2}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} \frac{1 - \alpha}{1 - \frac{1}{\gamma_u}} G_{\mathbf{a}}^{\text{sample}}(\{0\}).$$

The identity $\|\mathcal{P}_m \mathbf{a}\|^2 = \|\mathbf{a}\|^2 - \|\mathcal{P}_\perp \mathbf{a}\|^2$ is immediate. \square

Lemma .9 (Deterministic limit of $\mathbf{a}^\top \mathcal{P}_m \mathbf{b}$, $\mathbf{a}^\top \mathcal{P}_\perp \mathbf{b}$). *Consider the setup of Section 3, let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^p$ be fixed vectors and let \mathcal{P}_m denote the projection onto the top m principal components of \mathbf{S}_u , with $m/p \rightarrow \alpha$. Define the signed population spectral measure*

$$G_{\mathbf{a}, \mathbf{b}}(\tau) = \frac{1}{\|\mathbf{a}\| \|\mathbf{b}\|} \sum_j (\mathbf{a}^\top \mathbf{v}_j)(\mathbf{b}^\top \mathbf{v}_j) \mathbf{1}\{\lambda_j \leq \tau\}, \quad (34)$$

where $(\mathbf{v}_j, \lambda_j)$ are the population eigenpairs of Σ . Let $G_{\mathbf{a}, \mathbf{b}}^{\text{sample}}$ denote the corresponding deterministic limiting signed sample spectral measure onto the eigenpairs of \mathbf{S}_u . Then,

$$\mathbf{a}^\top \mathcal{P}_m \mathbf{b} = \mathbf{a}^\top \mathbf{b} - \mathbf{a}^\top \mathcal{P}_\perp \mathbf{b}, \quad (35)$$

and, almost surely,

$$\frac{\mathbf{a}^\top \mathcal{P}_\perp \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} \xrightarrow{a.s.} \begin{cases} \int_{\lambda_-}^{\lambda_t} g_{\mathbf{a}, \mathbf{b}}^{\text{sample}}(\lambda), & \gamma_u < 1, \\ G_{\mathbf{a}, \mathbf{b}}^{\text{sample}}(\{0\}) + \int_{\lambda_-}^{\lambda_t} g_{\mathbf{a}, \mathbf{b}}^{\text{sample}}(\lambda), & \gamma_u > 1, \alpha < 1/\gamma_u, \\ \frac{1 - \alpha}{1 - \frac{1}{\gamma_u}} G_{\mathbf{a}, \mathbf{b}}^{\text{sample}}(\{0\}), & \gamma_u > 1, \alpha \geq 1/\gamma_u. \end{cases} \quad (36)$$

Moreover, $G_{\mathbf{a}, \mathbf{b}}^{\text{sample}}(\{0\})$ and the bulk density of $g_{\mathbf{a}, \mathbf{b}}^{\text{sample}}$ are given by the formulas of Lemma .7 with $dG_{\mathbf{a}}$ replaced by the signed measure $dG_{\mathbf{a}, \mathbf{b}}$.

Proof. Define the signed mass density onto the sample eigenvectors of \mathbf{S}_u

$$\hat{G}_{\mathbf{a}, \mathbf{b}}^{\text{sample}}(\theta) := \frac{1}{\|\mathbf{a}\| \|\mathbf{b}\|} \sum_{i=1}^p (\mathbf{a}^\top \mathbf{u}_i)(\mathbf{b}^\top \mathbf{u}_i) \mathbf{1}\{\theta \leq \lambda_i(\mathbf{S}_u)\}, \quad (37)$$

so that its Stieltjes transform is

$$m_{\hat{G}_{\mathbf{a}, \mathbf{b}}^{\text{sample}}}(z) = \int \frac{1}{\lambda - z} d\hat{G}_{\mathbf{a}, \mathbf{b}}^{\text{sample}}(\lambda) = \mathbf{a}^\top (\mathbf{S}_u - zI)^{-1} \mathbf{b} = \text{Tr}[\mathbf{a} \mathbf{b}^\top (\mathbf{S}_u - zI)^{-1}]. \quad (38)$$

By Rubio & Mestre (2011, Theorem 1), for any deterministic matrix Θ with bounded trace norm (i.e., $\text{Tr}[(\Theta^\top \Theta)^{1/2}] = \|\Theta\|_{\text{tr}}$),

$$\text{Tr}[\Theta(\mathbf{S}_u - zI)^{-1}] - \delta_\Theta(z) \rightarrow 0 \quad \text{a.s.}$$

where $\delta_\Theta(z)$ is a deterministic function of z . Applying this with $\Theta = \mathbf{a}\mathbf{b}^\top$ and repeating the argument from (Green & Romanov, 2025) (in their paper, see Lemma 1 proven in Appendix A) yields the following almost sure weak convergence

$$\hat{G}_{\mathbf{a},\mathbf{b}}^{\text{sample}} \Rightarrow G_{\mathbf{a},\mathbf{b}}^{\text{sample}},$$

where the Stieltjes transform of $G_{\mathbf{a},\mathbf{b}}^{\text{sample}}$ is

$$m_{G_{\mathbf{a},\mathbf{b}}^{\text{sample}}}(z) = -\frac{1}{z} \int \frac{1}{1 + \tau \underline{m}_{\gamma_u}(z)} dG_{\mathbf{a},\mathbf{b}}(\tau). \quad (39)$$

Here, the companion Stieltjes transform is defined as $\underline{m}_{\gamma_u}(z) = \gamma_u m_{\gamma_u}(z) - (1 - \gamma_u) \frac{1}{z}$, and m_{γ_u} is a Stieltjes transform of a Marcenko-Pastur density f_{γ_u} . This is exactly the same expression as in Lemma 1 in (Green & Romanov, 2025) with the population measure $G_{\mathbf{a}}$ replaced by the signed measure $G_{\mathbf{a},\mathbf{b}}$. Hence, the atom at 0 and the bulk density of $G_{\mathbf{a},\mathbf{b}}^{\text{sample}}(\{0\})$ are obtained from the formulas of Lemma .7 by the same substitution. Now the mass of $\hat{G}_{\mathbf{a},\mathbf{b}}^{\text{sample}}$ on the discarded eigenspaces is

$$\mathbf{a}^\top \mathcal{P}_\perp \mathbf{b} = \sum_{i=m+1}^p (\mathbf{a}^\top \mathbf{u}_i)(\mathbf{b}^\top \mathbf{u}_i) \quad (40)$$

noting that the eigenvectors \mathbf{u}_i are sorted according to their eigenvalues, such that the set $\{m+1, \dots, p\}$ contains the indices of the $p-m$ smallest eigenvalues. Therefore, the same geometric argument as in Lemma .8 gives the three regimes above: discarded bulk only when $\gamma_u < 1$; bulk plus the full atom at 0 when $\gamma_u > 1$ and $\alpha < 1/\gamma_u$; and, when $\gamma_u > 1$ and $\alpha \geq 1/\gamma_u$, only the fraction

$$\frac{1 - \alpha}{1 - \frac{1}{\gamma_u}}$$

of the nullspace mass remains discarded, by the same random nullspace-completion argument. Finally,

$$\mathbf{a}^\top \mathcal{P}_m \mathbf{b} = \mathbf{a}^\top \mathbf{b} - \mathbf{a}^\top \mathcal{P}_\perp \mathbf{b},$$

which proves the result. \square

Lemma .10 (Deterministic limit of the effective projection $\|\Pi\mathbf{a}\|^2$). *Consider the setup of Section 3, let $\mathbf{X}_m = \mathbf{X}_l \mathbf{U}_m \in \mathbb{R}^{m \times m}$ and define*

$$\Pi_m := \mathbf{X}_m^\dagger \mathbf{X}_m \in \mathbb{R}^{m \times m}, \quad \Pi := \mathbf{U}_m \Pi_m \mathbf{U}_m^\top \in \mathbb{R}^{p \times p}.$$

Let the eigendecomposition of Σ_{mm} be $\Sigma_{mm} := \mathbf{U}_m^\top \Sigma \mathbf{U}_m$ with $(\tilde{\mathbf{v}}_i, \tilde{\lambda}_i)$ denoting the eigenpairs. For any fixed vector $\mathbf{a} \in \mathbb{R}^p$ with $\mathcal{P}_m \mathbf{a} \neq 0$, define

$$\tilde{\mathbf{a}} := \frac{\mathbf{U}_m^\top \mathbf{a}}{\|\mathcal{P}_m \mathbf{a}\|} \in \mathbb{R}^m,$$

and let $\tilde{G}_{\tilde{\mathbf{a}}}$ be the mass distribution of $\tilde{\mathbf{a}}$ onto the eigenvectors of Σ_{mm} :

$$\tilde{G}_{\tilde{\mathbf{a}}}(\tilde{\tau}) = \sum_i (\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}}_i)^2 \mathbf{1}\{\tilde{\lambda}_i \leq \tilde{\tau}\}.$$

Then, it holds almost surely

$$\tilde{G}_{\tilde{\mathbf{a}}}(\tau) \Rightarrow \left(1 - \frac{\bar{p}_{\mathbf{a},\mathbf{v}}^2}{\bar{p}_{\mathbf{a},\mathbf{a}} \bar{p}_{\mathbf{v},\mathbf{v}}}\right) \delta_1 + \frac{\bar{p}_{\mathbf{a},\mathbf{v}}^2}{\bar{p}_{\mathbf{a},\mathbf{a}} \bar{p}_{\mathbf{v},\mathbf{v}}} \delta_{\tilde{\lambda}},$$

and

$$\frac{\mathbf{a}^\top \Pi \mathbf{a}}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} \bar{p}_{\mathbf{a},\mathbf{a}} \cdot \begin{cases} 1, & \gamma_{\text{eff}} < 1, \\ 1 - \int \frac{\gamma_{\text{eff}} - 1}{\gamma_{\text{eff}} - 1 + \tilde{\tau}} d\tilde{G}_{\tilde{\mathbf{a}}}(\tilde{\tau}), & \gamma_{\text{eff}} > 1. \end{cases} \quad (41)$$

with $\bar{p}_{\mathbf{a},\mathbf{a}}, \bar{p}_{\mathbf{v},\mathbf{v}}, \bar{p}_{\mathbf{a},\mathbf{v}}$ defined via Equations (32) and (36).

1430 *Proof.* We start by computing explicitly $\tilde{G}_{\tilde{\mathbf{a}}}$. Recall that

$$1431 \quad \Sigma_{mm} = \mathbf{U}_m^\top \Sigma \mathbf{U}_m = \mathbf{I}_m + (\lambda - 1)(\mathbf{U}_m^\top \mathbf{v})(\mathbf{U}_m^\top \mathbf{v})^\top.$$

1433 Then, Σ_{mm} has $m - 1$ eigenvalues equal to 1 and the remaining one is

$$1435 \quad \tilde{\lambda} = 1 + (\lambda - 1)\|\mathcal{P}_m \mathbf{v}\|^2 \quad \Rightarrow \quad \bar{\lambda} = 1 + (\lambda - 1)\bar{p}_{\mathbf{v},\mathbf{v}},$$

1437 where the convergence occurs almost surely. The corresponding unit eigenvector is $\tilde{\mathbf{v}} = \frac{\mathbf{U}_m^\top \mathbf{v}}{\|\mathbf{U}_m^\top \mathbf{v}\|} = \frac{\mathbf{U}_m^\top \mathbf{v}}{\|\mathcal{P}_m \mathbf{v}\|}$. Therefore the

1438 spectral measure $\tilde{G}_{\tilde{\mathbf{a}}}$ of

$$1439 \quad \tilde{\mathbf{a}} = \frac{\mathbf{U}_m^\top \mathbf{a}}{\|\mathcal{P}_m \mathbf{a}\|}$$

1442 with respect to the eigenbasis of Σ_{mm} is supported on $\{1, \tilde{\lambda}\}$, and hence

$$1444 \quad \tilde{G}_{\tilde{\mathbf{a}}} = (1 - (\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}})^2)\delta_1 + (\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}})^2\delta_{\tilde{\lambda}},$$

1445 with

$$1446 \quad (\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}})^2 = \frac{((\mathbf{U}_m^\top \mathbf{a})^\top (\mathbf{U}_m^\top \mathbf{v}))^2}{\|\mathcal{P}_m \mathbf{a}\|^2 \|\mathcal{P}_m \mathbf{v}\|^2} = \frac{(\mathbf{a}^\top \mathbf{U}_m \mathbf{U}_m^\top \mathbf{v})^2}{\|\mathcal{P}_m \mathbf{a}\|^2 \|\mathcal{P}_m \mathbf{v}\|^2} = \frac{(\mathbf{a}^\top \mathcal{P}_m \mathbf{v})^2}{\|\mathcal{P}_m \mathbf{a}\|^2 \|\mathcal{P}_m \mathbf{v}\|^2}.$$

1449 Thus, $\tilde{G}_{\tilde{\mathbf{a}}}$ is a two-point measure, with the original atom at λ replaced by an atom at $\tilde{\lambda}$, and with weight determined by the

1450 projected overlap. Almost surely, it holds that:

$$1452 \quad \tilde{G}_{\tilde{\mathbf{a}}}(\tau) \Rightarrow \left(1 - \frac{\bar{p}_{\mathbf{a},\mathbf{v}}^2}{\bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{v},\mathbf{v}}}\right)\delta_1 + \frac{\bar{p}_{\mathbf{a},\mathbf{v}}^2}{\bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{v},\mathbf{v}}}\delta_{\tilde{\lambda}}.$$

1455 **Case $m < n_l$.** In this regime, $\mathbf{X}_m \in \mathbb{R}^{n_l \times m}$ has full column rank almost surely, so $\Pi_m = \mathbf{I}_m$. Hence $\tilde{\mathbf{a}}^\top \Pi_m \tilde{\mathbf{a}} = 1$,

1456 giving

$$1458 \quad \frac{\mathbf{a}^\top \Pi \mathbf{a}}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} \bar{p}_{\mathbf{a},\mathbf{a}},$$

1460 as $\frac{\|\mathcal{P}_m \mathbf{a}\|^2}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} \bar{p}_{\mathbf{a},\mathbf{a}}$ from Lemma .8.

1462 **Case $m > n_l$.** In this regime, \mathbf{X}_m has a nullspace of dimension $m - n_l$, and

$$1464 \quad \tilde{\mathbf{a}}^\top \Pi_m \tilde{\mathbf{a}} = 1 - \tilde{\mathbf{a}}^\top \Pi_{m,\text{null}} \tilde{\mathbf{a}}.$$

1466 Conditioned on \mathbf{U}_m , the rows of \mathbf{X}_m are Gaussian with covariance $\Sigma_{mm} = \mathbf{U}_m^\top \Sigma \mathbf{U}_m$. Therefore, by the nullspace mass

1467 formula from Lemma .7 applied to the m -dimensional model with aspect ratio $\gamma_{\text{eff}} = m/n_l > 1$,

$$1468 \quad \tilde{\mathbf{a}}^\top \Pi_{m,\text{null}} \tilde{\mathbf{a}} \xrightarrow{a.s.} \int \frac{\gamma_{\text{eff}} - 1}{\gamma_{\text{eff}} - 1 + \tilde{\tau}} d\tilde{G}_{\tilde{\mathbf{a}}}(\tilde{\tau}). \quad (42)$$

1471 Finally, noting that $\frac{\|\mathcal{P}_m \mathbf{a}\|^2}{\|\mathbf{a}\|^2} \xrightarrow{a.s.} \bar{p}_{\mathbf{a},\mathbf{a}}$ according to Lemma .8 and combining the two terms via the continuous mapping

1472 theorem gives the required result.

1473 □

1475 **Lemma .11** (Deterministic limit of the bilinear effective projection $\mathbf{a}^\top \Pi \mathbf{b}$). *Consider the setup of Section 3, let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^p$,*

1476 *with $\mathcal{P}_m \mathbf{a} \neq 0, \mathcal{P}_m \mathbf{b} \neq 0$, and define*

$$1478 \quad \tilde{\mathbf{a}} := \frac{\mathbf{U}_m^\top \mathbf{a}}{\|\mathcal{P}_m \mathbf{a}\|} \in \mathbb{R}^m, \quad \tilde{\mathbf{b}} := \frac{\mathbf{U}_m^\top \mathbf{b}}{\|\mathcal{P}_m \mathbf{b}\|} \in \mathbb{R}^m.$$

1481 Let $\tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}}$ be the signed joint mass distribution of $\tilde{\mathbf{a}}, \tilde{\mathbf{b}}$ onto the eigenvectors $\tilde{\mathbf{v}}_i$ of Σ_{mm} :

$$1483 \quad \tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}}(\tilde{\tau}) = \sum_i (\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}}_i)(\tilde{\mathbf{b}}^\top \tilde{\mathbf{v}}_i) \mathbf{1}\{\tilde{\lambda}_i \leq \tilde{\tau}\}.$$

1484

Then, it holds almost surely

$$\tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}}(\tau) \Rightarrow \left(\frac{\bar{p}_{\mathbf{a},\mathbf{b}}}{\bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{b},\mathbf{b}}} - \frac{\bar{p}_{\mathbf{a},\mathbf{v}}\bar{p}_{\mathbf{b},\mathbf{v}}}{\bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{b},\mathbf{b}}\bar{p}_{\mathbf{v},\mathbf{v}}} \right) \delta_1 + \frac{\bar{p}_{\mathbf{a},\mathbf{v}}\bar{p}_{\mathbf{b},\mathbf{v}}}{\bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{b},\mathbf{b}}\bar{p}_{\mathbf{v},\mathbf{v}}} \delta_{\tilde{\lambda}},$$

and

$$\frac{\mathbf{a}^\top \Pi \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} \xrightarrow{a.s.} \begin{cases} \bar{p}_{\mathbf{a},\mathbf{b}}, & \gamma_{\text{eff}} < 1, \\ \bar{p}_{\mathbf{a},\mathbf{b}} - \bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{b},\mathbf{b}} \int \frac{\gamma_{\text{eff}} - 1}{\gamma_{\text{eff}} - 1 + \tilde{\tau}} d\tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}}(\tilde{\tau}), & \gamma_{\text{eff}} > 1. \end{cases} \quad (43)$$

with $\bar{p}_{\mathbf{a},\mathbf{a}}, \bar{p}_{\mathbf{b},\mathbf{b}}, \bar{p}_{\mathbf{v},\mathbf{v}}, \bar{p}_{\mathbf{a},\mathbf{v}}, \bar{p}_{\mathbf{b},\mathbf{v}}$ defined via Eq. (32) and (36).

Proof. The proof follows the steps of the proof of Lemma .10. The mass distribution $\tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}}$ can be written as

$$\tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}} = (\tilde{\mathbf{a}}^\top \tilde{\mathbf{b}} - (\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}})(\tilde{\mathbf{b}}^\top \tilde{\mathbf{v}}))\delta_1 + (\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}})(\tilde{\mathbf{b}}^\top \tilde{\mathbf{v}})\delta_{\tilde{\lambda}},$$

with

$$(\tilde{\mathbf{a}}^\top \tilde{\mathbf{v}})(\tilde{\mathbf{b}}^\top \tilde{\mathbf{v}}) = \frac{((\mathbf{U}_m^\top \mathbf{a})^\top (\mathbf{U}_m^\top \mathbf{v}))((\mathbf{U}_m^\top \mathbf{b})^\top (\mathbf{U}_m^\top \mathbf{v}))}{\|\mathcal{P}_m \mathbf{a}\| \|\mathcal{P}_m \mathbf{b}\| \|\mathcal{P}_m \mathbf{v}\|^2} = \frac{(\mathbf{a}^\top \mathcal{P}_m \mathbf{v})(\mathbf{b}^\top \mathcal{P}_m \mathbf{v})}{\|\mathcal{P}_m \mathbf{a}\| \|\mathcal{P}_m \mathbf{b}\| \|\mathcal{P}_m \mathbf{v}\|^2},$$

which converges almost surely to $\frac{\bar{p}_{\mathbf{a},\mathbf{v}}\bar{p}_{\mathbf{b},\mathbf{v}}}{\bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{b},\mathbf{b}}\bar{p}_{\mathbf{v},\mathbf{v}}}$ and

$$\tilde{\mathbf{a}}^\top \tilde{\mathbf{b}} = \frac{(\mathbf{U}_m^\top \mathbf{a})^\top (\mathbf{U}_m^\top \mathbf{b})}{\|\mathcal{P}_m \mathbf{a}\| \|\mathcal{P}_m \mathbf{b}\|} \xrightarrow{a.s.} \frac{\bar{p}_{\mathbf{a},\mathbf{b}}}{\bar{p}_{\mathbf{a},\mathbf{a}}\bar{p}_{\mathbf{b},\mathbf{b}}}$$

This readily gives the almost sure limit of $\tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}}(\tau)$. Next, observe

$$\frac{\mathbf{a}^\top \Pi \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} = \frac{1}{\|\mathbf{a}\| \|\mathbf{b}\|} (\mathbf{U}_m^\top \mathbf{a})^\top \Pi_m (\mathbf{U}_m^\top \mathbf{b}) = \frac{\|\mathcal{P}_m \mathbf{a}\| \|\mathcal{P}_m \mathbf{b}\|}{\|\mathbf{a}\| \|\mathbf{b}\|} \tilde{\mathbf{a}}^\top \Pi_m \tilde{\mathbf{b}}. \quad (44)$$

Combining this with the almost sure limits from Lemma .8, with the same steps as in the proof of Lemma .10 and with the mass distribution of $\tilde{G}_{\tilde{\mathbf{a}},\tilde{\mathbf{b}}}$ proves the lemma. \square

A.7. PHASE TRANSITION FOR $\gamma_l > 1$

Here, we derive parametric equations for the two phase transitions we observe in the phase diagrams for α in Figure 3. For small downstream tasks, $\gamma_l > 1$, there is a sharp phase transition due to competing local minima at different sides of the double descent (local minimum at large vs small α), which we tackle first. We find a proxy of the parametric curve of this phase transition by comparing the values of the error at the endpoints in α , and check when $E_\infty^{\text{gen}}(\alpha \rightarrow 0) = E_\infty^{\text{gen}}(\alpha = 1)$. Note that the left endpoint is not exactly 0, as we always retain the first principal component of the data, but in the high-dimensional regime $p \rightarrow \infty$, α approaches 0.

We define signed normalized overlaps as

$$\rho_{\mathbf{a},\mathbf{b}} := \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}, \quad \eta_{\mathbf{a},\mathbf{b}} := \rho_{\mathbf{a},\mathbf{b}}^2.$$

In particular, $\eta_{\mathbf{a},\mathbf{v}}$ denotes the squared population-spike overlap of a fixed vector \mathbf{a} with the spike eigenvector \mathbf{v} . For bilinear quantities, the signs of the corresponding ρ 's must be retained.

Define the limiting value of the population and sample spike eigenvector overlap $(\mathbf{v}^\top \mathbf{u}_1)^2 \rightarrow c$, which depends on whether the spike is below or above the BBP transition (Green & Romanov, 2025):

$$c := \begin{cases} \frac{1 - \frac{\gamma_u}{(\lambda-1)^2}}{1 + \frac{\gamma_u}{\lambda-1}} & \gamma_u < (\lambda-1)^2 \\ 0 & \gamma_u \geq (\lambda-1)^2 \end{cases}.$$

Endpoint values for the eigenvalues and projections. At the left endpoint, meaning $m = 1$ and hence $\alpha \rightarrow 0$, the retained direction is the empirical spike direction. Therefore, the limiting value of projections for any fixed vectors \mathbf{a}, \mathbf{b} are given by

$$\begin{aligned}\bar{p}_{\mathbf{a},\mathbf{a}}(0) &= \eta_{\mathbf{a},\mathbf{v}c}, & \bar{p}_{\perp,\mathbf{a},\mathbf{a}}(0) &= 1 - \eta_{\mathbf{a},\mathbf{v}c}, \\ \bar{p}_{\mathbf{a},\mathbf{b}}(0) &= \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}c}, & \bar{p}_{\perp,\mathbf{a},\mathbf{b}}(0) &= \rho_{\mathbf{a},\mathbf{b}} - \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}c}.\end{aligned}$$

At the right endpoint, $\alpha = 1$, the projection is the identity, $\mathcal{P}_m = \mathbf{I}_p$. Hence, we have

$$\begin{aligned}\bar{p}_{\mathbf{a},\mathbf{a}}(1) &= 1, & \bar{p}_{\perp,\mathbf{a},\mathbf{a}}(1) &= 0 \\ \bar{p}_{\mathbf{a},\mathbf{b}}(1) &= \rho_{\mathbf{a},\mathbf{b}}, & \bar{p}_{\perp,\mathbf{a},\mathbf{b}}(1) &= 0.\end{aligned}$$

Next, let us analyse the effective limiting projections $\bar{\pi}$. At $\alpha = 0$, the effective aspect ratio satisfies $\gamma_{\text{eff}} = 0$. Hence the compressed task design is overdetermined, and the task projection is exact inside the retained subspace:

$$\begin{aligned}\bar{\pi}_{\mathbf{a},\mathbf{a}}(0) &= \bar{p}_{\mathbf{a},\mathbf{a}}(0) = \eta_{\mathbf{a},\mathbf{v}c}, \\ \bar{\pi}_{\mathbf{a},\mathbf{b}}(0) &= \bar{p}_{\mathbf{a},\mathbf{b}}(0) = \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}c}.\end{aligned}$$

At $\alpha = 1$, we have $\bar{\lambda} = \lambda$ and $\gamma_{\text{eff}} = \gamma_l$. Therefore

$$\begin{aligned}\bar{\pi}_{\mathbf{a},\mathbf{a}}(1) &= \begin{cases} 1, & \gamma_l < 1, \\ \eta_{\mathbf{a},\mathbf{v}} \frac{\lambda}{\gamma_l - 1 + \lambda} + (1 - \eta_{\mathbf{a},\mathbf{v}}) \frac{1}{\gamma_l}, & \gamma_l > 1. \end{cases} \\ \bar{\pi}_{\mathbf{a},\mathbf{b}}(1) &= \begin{cases} \rho_{\mathbf{a},\mathbf{b}}, & \gamma_l < 1, \\ \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}} \frac{\lambda}{\gamma_l - 1 + \lambda} + (\rho_{\mathbf{a},\mathbf{b}} - \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}}) \frac{1}{\gamma_l}, & \gamma_l > 1. \end{cases}\end{aligned}$$

Furthermore, note that at $\alpha = 1$ we have $\bar{\lambda} = \lambda$, and for $\alpha = 0$ define:

$$\bar{\lambda}_0 = 1 + (\lambda - 1)c.$$

Endpoint values for the generalisation error. At $\alpha = 1$, we have $\mathcal{P}_m = \mathbf{I}_p$, hence

$$\bar{p}_{\perp,\mathbf{w}^*,\mathbf{w}^*} = 0, \quad \bar{p}_{\perp,\mathbf{w}^*,\mathbf{v}} = 0, \quad \bar{\sigma}_{\text{eff}}^2 = 0.$$

Therefore,

$$E_{\infty}^{\text{gen}}(1) = \bar{w}^* \left(1 - \bar{\pi}_{\mathbf{w}^*,\mathbf{w}^*}(1) \right) + \frac{\sigma^2}{|\gamma_l - 1|} + \bar{w}^* \sum_{i=1}^T (\nu_i - 1) \left(\rho_{\mathbf{w}^*,\mathbf{v}_{\text{new},i}} - \bar{\pi}_{\mathbf{w}^*,\mathbf{v}_{\text{new},i}}(1) \right)^2 + \sigma^2,$$

where for the relevant regime for this phase transition ($\gamma_l > 1$)

$$\begin{aligned}\bar{\pi}_{\mathbf{w}^*,\mathbf{w}^*}(1) &= \eta_{\mathbf{w}^*,\mathbf{v}} \frac{\lambda}{\gamma_l - 1 + \lambda} + (1 - \eta_{\mathbf{w}^*,\mathbf{v}}) \frac{1}{\gamma_l}, \\ \bar{\pi}_{\mathbf{w}^*,\mathbf{v}_{\text{new},i}}(1) &= \rho_{\mathbf{w}^*,\mathbf{v}}\rho_{\mathbf{v}_{\text{new},i},\mathbf{v}} \frac{\lambda}{\gamma_l - 1 + \lambda} + (\rho_{\mathbf{w}^*,\mathbf{v}_{\text{new},i}} - \rho_{\mathbf{w}^*,\mathbf{v}}\rho_{\mathbf{v}_{\text{new},i},\mathbf{v}}) \frac{1}{\gamma_l}.\end{aligned}$$

At $\alpha = 0$ ($\gamma_{\text{eff}} = 0$), the variance term in E_{∞}^{est} vanishes. Moreover,

$$\bar{\pi}_{\mathbf{w}^*,\mathbf{w}^*}(0) = \eta_{\mathbf{w}^*,\mathbf{v}c}, \quad \bar{\pi}_{\mathbf{v},\mathbf{v}}(0) = c, \quad \bar{p}_{\perp,\mathbf{w}^*,\mathbf{v}}(0) = \rho_{\mathbf{w}^*,\mathbf{v}}(1 - c).$$

Thus, we have

$$\begin{aligned}E_{\infty}^{\text{gen}}(0) &= \bar{w}^* \left[1 - \eta_{\mathbf{w}^*,\mathbf{v}c} + \frac{(\lambda - 1)^2}{\bar{\lambda}_0^2} \eta_{\mathbf{w}^*,\mathbf{v}}(1 - c)^2 c \right] \\ &+ \bar{w}^* \sum_{i=1}^T (\nu_i - 1) \left[\rho_{\mathbf{w}^*,\mathbf{v}_{\text{new},i}} - \rho_{\mathbf{w}^*,\mathbf{v}}\rho_{\mathbf{v}_{\text{new},i},\mathbf{v}}c - \frac{\lambda - 1}{\bar{\lambda}_0} \rho_{\mathbf{w}^*,\mathbf{v}}(1 - c)\rho_{\mathbf{v}_{\text{new},i},\mathbf{v}}c \right]^2 + \sigma^2.\end{aligned}$$

Noting that $\bar{\lambda}_0 = 1 + (\lambda - 1)c$, this simplifies further:

$$E_\infty^{\text{gen}}(0) = \bar{w}^* \left[1 - \eta_{\mathbf{w}^*, \mathbf{v}} c \frac{\lambda}{\lambda_0} \left(2 - \frac{\lambda}{\lambda_0} \right) + \sum_{i=1}^T (\nu_i - 1) \left(\rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \rho_{\mathbf{w}^*, \mathbf{v}} \rho_{\mathbf{v}, \mathbf{v}_{\text{new}, i}} c \frac{\lambda}{\lambda_0} \right)^2 \right] + \sigma^2.$$

Let $S := \frac{\bar{w}^*}{\sigma^2}$, and define

$$r_i := \rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}}, \quad q_i := \rho_{\mathbf{w}^*, \mathbf{v}} \rho_{\mathbf{v}, \mathbf{v}_{\text{new}, i}}, \quad \eta := \eta_{\mathbf{w}^*, \mathbf{v}}.$$

Then the normalized $\alpha = 0$ endpoint is

$$B_0(c) := 1 - \eta c \frac{\lambda}{\lambda_0} \left(2 - \frac{\lambda}{\lambda_0} \right) + \sum_{i=1}^T (\nu_i - 1) \left(r_i - q_i \frac{\lambda c}{\lambda_0} \right)^2.$$

Thus

$$\frac{E_\infty^{\text{gen}}(0)}{\bar{w}^*} = B_0(c) + \frac{1}{S}.$$

For $\alpha = 1$ and $\gamma_l > 1$, define

$$A(\gamma_l) := \frac{\lambda}{\gamma_l - 1 + \lambda}, \quad B(\gamma_l) := \frac{1}{\gamma_l}.$$

Then

$$B_1(\gamma_l) := 1 - [\eta A(\gamma_l) + (1 - \eta)B(\gamma_l)] + \sum_{i=1}^T (\nu_i - 1) [r_i - q_i A(\gamma_l) - (r_i - q_i)B(\gamma_l)]^2.$$

The phase boundary in the (γ_u, γ_l) -plane is therefore

$$\boxed{B_0(c(\gamma_u)) = B_1(\gamma_l) + \frac{1}{S(\gamma_l - 1)}} \quad (\gamma_l > 1).$$

Specialization to $\Sigma_{\text{new}} = \Sigma$. In this case, $T = 1$, $\nu_1 = \lambda$, and

$$r_1 = q_1 = \rho_{\mathbf{w}^*, \mathbf{v}}, \quad \eta = \rho_{\mathbf{w}^*, \mathbf{v}}^2.$$

Hence

$$\begin{aligned} B_0(c) &= 1 + \frac{\eta}{\bar{\lambda}_0^2} (c\lambda^2 - 2c\lambda\bar{\lambda}_0 + (\lambda - 1)(\bar{\lambda}_0^2 - 2\bar{\lambda}_0\lambda c + \lambda^2 c^2)) \\ &= 1 + \eta(\lambda - 1) - \frac{\eta\lambda^2 c}{1 + (\lambda - 1)c}, \end{aligned}$$

while

$$B_1 = (1 - \eta) \frac{\gamma_l - 1}{\gamma_l} + \eta \frac{\lambda\gamma_l(\gamma_l - 1)}{(\gamma_l - 1 + \lambda)^2}.$$

Thus, the phase transition is:

$$1 + \eta(\lambda - 1) - \frac{\eta\lambda^2 c}{1 + (\lambda - 1)c} = (1 - \eta) \frac{\gamma_l - 1}{\gamma_l} + \eta \frac{\lambda\gamma_l(\gamma_l - 1)}{(\gamma_l - 1 + \lambda)^2} + \frac{1}{S(\gamma_l - 1)}$$

Specialization to $\Sigma_{\text{new}} = \Sigma$ and $\mathbf{w}^* = \sqrt{\bar{w}^*} \mathbf{v}$. Now everything simplifies even further. We have

$$B_0 = \frac{\lambda(1 - c)}{1 + (\lambda - 1)c},$$

and

$$B_1 = \frac{\lambda\gamma_l(\gamma_l - 1)}{(\gamma_l - 1 + \lambda)^2}.$$

Due to the two regimes of c , the proxy transition splits into two components:

$$\frac{\lambda\gamma_l(\gamma_l - 1)}{(\gamma_l - 1 + \lambda)^2} + \frac{1}{S(\gamma_l - 1)} = \begin{cases} \lambda & \lambda < 1 + \sqrt{\gamma_u} \\ \frac{\lambda}{(\lambda - 1)^2} \gamma_u & \lambda \geq 1 + \sqrt{\gamma_u} \end{cases}$$

A.8. PHASE TRANSITION FOR $\gamma_l < 1$

Next, we study the regime of larger downstream tasks $\gamma_l < 1$, and find the phase transition between the phases $\alpha = 1$ and $\alpha < 1$ by finding a parametric curve that describes the loss of stability of $\alpha = 1$, i.e., when $\frac{d}{d\alpha} E_\infty^{\text{gen}}|_{\alpha=1} = 0$.

Derivatives of the pretraining projection at $\alpha = 1$. Let $\lambda_\pm = (1 \pm \sqrt{\gamma_u})^2$, and define

$$H_{\mathbf{a}}(\theta) := (1 - \eta_{\mathbf{a},\mathbf{v}}) + \eta_{\mathbf{a},\mathbf{v}} \frac{\lambda \gamma_u}{\lambda^2 - \lambda(1 - \gamma_u + \theta) + \theta}.$$

Since

$$g_{\mathbf{a}}^{\text{sample}}(\theta) = H_{\mathbf{a}}(\theta) f_{\gamma_u}(\theta),$$

the derivative of $\bar{p}_{\perp, \mathbf{a}, \mathbf{a}}$ at $\alpha = 1$ is

$$\left. \frac{d}{d\alpha} \bar{p}_{\perp, \mathbf{a}, \mathbf{a}}(\alpha) \right|_{\alpha=1^-} = \begin{cases} -H_{\mathbf{a}}(\lambda_-), & \gamma_u < 1, \\ -\frac{G_{\mathbf{a}}^{\text{sample}}(\{0\})}{1 - \frac{1}{\gamma_u}}, & \gamma_u > 1. \end{cases}$$

Equivalently,

$$\left. \frac{d}{d\alpha} \bar{p}_{\perp, \mathbf{a}, \mathbf{a}}(\alpha) \right|_{\alpha=1^-} = \begin{cases} -\left[(1 - \eta_{\mathbf{a},\mathbf{v}}) + \eta_{\mathbf{a},\mathbf{v}} \frac{\lambda \gamma_u}{\lambda^2 - \lambda(1 - \gamma_u + \lambda_-) + \lambda_-} \right], & \gamma_u < 1, \\ -\left[(1 - \eta_{\mathbf{a},\mathbf{v}}) + \eta_{\mathbf{a},\mathbf{v}} \frac{\gamma_u}{\gamma_u - 1 + \lambda} \right], & \gamma_u > 1. \end{cases}$$

Therefore

$$\left. \frac{d}{d\alpha} \bar{p}_{\mathbf{a}, \mathbf{a}}(\alpha) \right|_{\alpha=1^-} = - \left. \frac{d}{d\alpha} \bar{p}_{\perp, \mathbf{a}, \mathbf{a}}(\alpha) \right|_{\alpha=1^-}.$$

For signed bilinear projections, replace

$$\eta_{\mathbf{a},\mathbf{v}} \text{ by } \rho_{\mathbf{a},\mathbf{v}} \rho_{\mathbf{b},\mathbf{v}},$$

and replace the constant bulk contribution $1 - \eta_{\mathbf{a},\mathbf{v}}$ by

$$\rho_{\mathbf{a},\mathbf{b}} - \rho_{\mathbf{a},\mathbf{v}} \rho_{\mathbf{b},\mathbf{v}}.$$

Derivative of the effective projection at $\alpha = 1$. For $\gamma_l < 1$, the compressed task design remains overdetermined in a neighbourhood of $\alpha = 1$. Thus

$$\bar{\pi}_{\mathbf{a}, \mathbf{a}}(\alpha) = \bar{p}_{\mathbf{a}, \mathbf{a}}(\alpha), \quad \gamma_l < 1,$$

and hence

$$\left. \frac{d}{d\alpha} \bar{\pi}_{\mathbf{a}, \mathbf{a}}(\alpha) \right|_{\alpha=1^-} = \left. \frac{d}{d\alpha} \bar{p}_{\mathbf{a}, \mathbf{a}}(\alpha) \right|_{\alpha=1^-}.$$

Stability of $\alpha = 1$ for $\gamma_l < 1$. Assume $\gamma_l < 1$ and let $\alpha \rightarrow 1^-$. Define

$$D_{\mathbf{w}^*} := \left. \frac{d}{d\alpha} \bar{p}_{\mathbf{w}^*, \mathbf{w}^*}(\alpha) \right|_{\alpha=1^-}.$$

At $\alpha = 1$, $\bar{p}_{\perp, \mathbf{a}, \mathbf{b}} = 0$ and $\sigma_{\text{eff}}^2 = 0$, thus simplifying

$$\left. \frac{d}{d\alpha} E_\infty^{\text{gen}}(\alpha) \right|_{\alpha=1^-} = \frac{\gamma_l \sigma^2}{(1 - \gamma_l)^2} - \frac{\bar{w}^*}{1 - \gamma_l} D_{\mathbf{w}^*}.$$

Consequently, $\alpha = 1$ is locally stable iff

$$\left. \frac{d}{d\alpha} E_\infty^{\text{gen}}(\alpha) \right|_{\alpha=1^-} \leq 0,$$

or equivalently

$$\gamma_l \leq \frac{SD_{\mathbf{w}^*}}{1 + SD_{\mathbf{w}^*}}, \quad S := \frac{\bar{w}^*}{\sigma^2}.$$

The equality gives the local stability boundary of the full representation:

$$D_{\mathbf{w}^*} = \begin{cases} (1 - \eta) + \eta \frac{\lambda \gamma_u}{(\lambda - 1 + \sqrt{\gamma_u})^2}, & \gamma_u < 1, \\ (1 - \eta) + \eta \frac{\gamma_u}{\gamma_u - 1 + \lambda}, & \gamma_u > 1, \end{cases}$$

$$\gamma_l^{\text{crit}} = \frac{S[(1 - \eta) + \eta D_{\mathbf{v}}]}{1 + S[(1 - \eta) + \eta D_{\mathbf{v}}]},$$

where

$$D_{\mathbf{v}} = \begin{cases} \frac{\lambda \gamma_u}{(\lambda - 1 + \sqrt{\gamma_u})^2}, & \gamma_u < 1, \\ \frac{\gamma_u}{\gamma_u - 1 + \lambda}, & \gamma_u > 1. \end{cases}$$

A.9. INFINITE DATA LIMITS

We consider what happens to the errors and optimal α when $\gamma_u \rightarrow 0$ or $\gamma_l \rightarrow 0$.

Prior data limit $\gamma_u \rightarrow 0$. In this limit, the empirical spike aligns with the population spike, and the retained subspace contains the spike direction together with an α -fraction of the isotropic bulk.

For any fixed vectors \mathbf{a}, \mathbf{b} , we have

$$\begin{aligned} \bar{p}_{\mathbf{a},\mathbf{a}} &\rightarrow \eta_{\mathbf{a},\mathbf{v}} + (1 - \eta_{\mathbf{a},\mathbf{v}})\alpha, \\ \bar{p}_{\perp,\mathbf{a},\mathbf{a}} &\rightarrow (1 - \eta_{\mathbf{a},\mathbf{v}})(1 - \alpha), \\ \bar{p}_{\mathbf{a},\mathbf{b}} &\rightarrow \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}} + \alpha(\rho_{\mathbf{a},\mathbf{b}} - \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}}), \\ \bar{p}_{\perp,\mathbf{a},\mathbf{b}} &\rightarrow (1 - \alpha)(\rho_{\mathbf{a},\mathbf{b}} - \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}}). \end{aligned}$$

In particular,

$$\bar{p}_{\mathbf{v},\mathbf{v}} \rightarrow 1, \quad \bar{p}_{\perp,\mathbf{v},\mathbf{v}} \rightarrow 0, \quad \bar{\lambda} \rightarrow \lambda.$$

For the effective projections $\bar{\pi}$, the behaviour depends on γ_{eff} .

If $\gamma_{\text{eff}} < 1$, then

$$\bar{\pi}_{\mathbf{a},\mathbf{a}} \rightarrow \bar{p}_{\mathbf{a},\mathbf{a}}, \quad \bar{\pi}_{\mathbf{a},\mathbf{b}} \rightarrow \bar{p}_{\mathbf{a},\mathbf{b}}.$$

If $\gamma_{\text{eff}} > 1$, define

$$A_\lambda := \frac{\lambda}{\gamma_{\text{eff}} - 1 + \lambda}, \quad A_1 := \frac{1}{\gamma_{\text{eff}}},$$

then

$$\begin{aligned} \bar{\pi}_{\mathbf{a},\mathbf{a}} &\rightarrow \eta_{\mathbf{a},\mathbf{v}}A_\lambda + (1 - \eta_{\mathbf{a},\mathbf{v}})\alpha A_1, \\ \bar{\pi}_{\mathbf{a},\mathbf{b}} &\rightarrow \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}}A_\lambda + \alpha(\rho_{\mathbf{a},\mathbf{b}} - \rho_{\mathbf{a},\mathbf{v}}\rho_{\mathbf{b},\mathbf{v}})A_1. \end{aligned}$$

Errors in the limit $\gamma_u \rightarrow 0$. In the infinite-prior-data limit, the empirical spike is recovered exactly. Hence

$$\bar{p}_{\mathbf{v},\mathbf{v}} \rightarrow 1, \quad \bar{p}_{\perp,\mathbf{w}^*,\mathbf{v}} \rightarrow 0, \quad \bar{\lambda} \rightarrow \lambda,$$

and

$$\bar{p}_{\mathbf{w}^*,\mathbf{w}^*} \rightarrow \eta + (1 - \eta)\alpha, \quad \bar{p}_{\perp,\mathbf{w}^*,\mathbf{w}^*} \rightarrow (1 - \eta)(1 - \alpha).$$

1760 Therefore

$$1761 \bar{\sigma}_{\text{eff}}^2 \rightarrow (1 - \eta)(1 - \alpha).$$

1762 Since the leak term vanishes,

$$1763 E_{\infty}^{\text{est}} \rightarrow \bar{w}^* (1 - \bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*}) + \frac{\min\{\gamma_{\text{eff}}, 1\}}{|\gamma_{\text{eff}} - 1|} [\bar{w}^* (1 - \eta)(1 - \alpha) + \sigma^2],$$

1764 where

$$1765 \bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*} = \begin{cases} \eta + (1 - \eta)\alpha, & \gamma_{\text{eff}} < 1, \\ \eta \frac{\lambda}{\gamma_{\text{eff}} - 1 + \lambda} + (1 - \eta) \frac{1}{\gamma_l}, & \gamma_{\text{eff}} > 1. \end{cases}$$

1766 Moreover, taking $\Sigma_{\text{new}} = \Sigma$, we obtain

$$1767 E_{\infty}^{\text{gen}} \rightarrow E_{\infty}^{\text{est}} + \bar{w}^* (\lambda - 1) \eta \left(\frac{\gamma_{\text{eff}} - 1}{\gamma_{\text{eff}} - 1 + \lambda} \right)^2 + \sigma^2.$$

1768 Finally, the training error satisfies

$$1769 E_{\infty}^{\text{train}} \rightarrow [\sigma^2 + \bar{w}^* (1 - \eta)(1 - \alpha)] \max\{0, 1 - \gamma_{\text{eff}}\}.$$

1770 **Downstream-data limit** $\gamma_l \rightarrow 0$. For any fixed $\alpha > 0$, we have $\gamma_{\text{eff}} = \alpha\gamma_l \rightarrow 0$, hence

$$1771 \bar{\pi}_{\mathbf{a}, \mathbf{a}} \rightarrow \bar{p}_{\mathbf{a}, \mathbf{a}}, \quad \bar{\pi}_{\mathbf{a}, \mathbf{b}} \rightarrow \bar{p}_{\mathbf{a}, \mathbf{b}}.$$

1772 Moreover, the variance factor satisfies

$$1773 \frac{\min\{\gamma_{\text{eff}}, 1\}}{|\gamma_{\text{eff}} - 1|} = \frac{\alpha\gamma_l}{1 - \alpha\gamma_l} \rightarrow 0,$$

1774 so all downstream finite-sample variance contributions vanish.

1775 The pretraining projections \bar{p} remain determined by γ_u .

1776 At $\alpha = 1$,

$$1777 \bar{p}_{\mathbf{a}, \mathbf{a}} = 1, \quad \bar{p}_{\perp, \mathbf{a}, \mathbf{a}} = 0, \quad \bar{\pi}_{\mathbf{a}, \mathbf{a}} \rightarrow 1,$$

1778 and

$$1779 \bar{p}_{\mathbf{a}, \mathbf{b}} = \rho_{\mathbf{a}, \mathbf{b}}, \quad \bar{p}_{\perp, \mathbf{a}, \mathbf{b}} = 0, \quad \bar{\pi}_{\mathbf{a}, \mathbf{b}} \rightarrow \rho_{\mathbf{a}, \mathbf{b}}.$$

1780 At the spike-retaining endpoint $\alpha \rightarrow 0$,

$$1781 \bar{\pi}_{\mathbf{a}, \mathbf{a}}(0) \rightarrow \bar{p}_{\mathbf{a}, \mathbf{a}}(0), \quad \bar{\pi}_{\mathbf{a}, \mathbf{b}}(0) \rightarrow \bar{p}_{\mathbf{a}, \mathbf{b}}(0),$$

1782 with \bar{p} determined by the prior overlap $c(\gamma_u)$.

1783 **Errors in the limit** $\gamma_l \rightarrow 0$. For fixed $\alpha > 0$, $\gamma_{\text{eff}} = \alpha\gamma_l \rightarrow 0$. Hence

$$1784 \bar{\pi}_{\mathbf{a}, \mathbf{a}} \rightarrow \bar{p}_{\mathbf{a}, \mathbf{a}}, \quad \bar{\pi}_{\mathbf{a}, \mathbf{b}} \rightarrow \bar{p}_{\mathbf{a}, \mathbf{b}},$$

1785 and

$$1786 \frac{\min\{\gamma_{\text{eff}}, 1\}}{|\gamma_{\text{eff}} - 1|} \rightarrow 0.$$

1787 Therefore the estimation error satisfies

$$1788 E_{\infty}^{\text{est}} \rightarrow \bar{w}^* \left[1 - \bar{p}_{\mathbf{w}^*, \mathbf{w}^*} + \frac{(\lambda - 1)^2}{\lambda^2} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}^2 \bar{p}_{\mathbf{v}, \mathbf{v}} \right].$$

The generalisation error satisfies

$$E_{\infty}^{\text{gen}} \rightarrow E_{\infty}^{\text{est}} + \bar{w}^* \sum_{i=1}^T (\nu_i - 1) \left(\rho_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \bar{p}_{\mathbf{w}^*, \mathbf{v}_{\text{new}, i}} - \frac{\lambda - 1}{\lambda} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}} \bar{p}_{\mathbf{v}, \mathbf{v}_{\text{new}, i}} \right)^2 + \sigma^2. \quad (45)$$

The training error satisfies

$$E_{\infty}^{\text{train}} \rightarrow \sigma^2 + \bar{w}^* \left[\bar{p}_{\perp, \mathbf{w}^*, \mathbf{w}^*} + \frac{\lambda - 1}{\lambda} \bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}^2 \right]. \quad (46)$$

Specialization to $\Sigma_{\text{new}} = \Sigma$. We further specialize the result in the limit to the setting where the validation data is the same as pretraining and downstream data.

Corollary .12. Assume $\gamma_u = p/n_u \rightarrow 0$, so that the empirical spike direction converges to the population spike \mathbf{v} . Let $\eta = \eta_{\mathbf{w}^*, \mathbf{v}}$, $\gamma_{\text{eff}} = \alpha\gamma_l$, $S = \bar{w}^*/\sigma^2$. Then the estimation error becomes

$$E_{\infty}^{\text{est}} \rightarrow \bar{w}^*(1 - \bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*}) + \frac{\min\{\gamma_{\text{eff}}, 1\}}{|\gamma_{\text{eff}} - 1|} [\bar{w}^*(1 - \eta)(1 - \alpha) + \sigma^2],$$

where

$$\bar{\pi}_{\mathbf{w}^*, \mathbf{w}^*} = \begin{cases} \eta + (1 - \eta)\alpha, & \gamma_{\text{eff}} < 1, \\ \eta \frac{\lambda}{\gamma_{\text{eff}} - 1 + \lambda} + (1 - \eta) \frac{\alpha}{\gamma_{\text{eff}}}, & \gamma_{\text{eff}} > 1. \end{cases}$$

For matched test covariance $\Sigma_{\text{new}} = \Sigma$,

$$E_{\infty}^{\text{gen}} \rightarrow E_{\infty}^{\text{est}} + \bar{w}^*(\lambda - 1)\eta \begin{cases} 0, & \gamma_{\text{eff}} < 1, \\ \left(\frac{\gamma_{\text{eff}} - 1}{\gamma_{\text{eff}} - 1 + \lambda} \right)^2, & \gamma_{\text{eff}} > 1, \end{cases} + \sigma^2.$$

Finally,

$$E_{\infty}^{\text{train}} \rightarrow [\sigma^2 + \bar{w}^*(1 - \eta)(1 - \alpha)] \max\{0, 1 - \gamma_{\text{eff}}\}.$$

Consequently, the phase transitions in this limit can either be derived by taking derivatives and limiting values α , or by substituting the limits directly in Corollaries 4.4 and 4.5. In this regime, the phase boundaries become independent of γ_u and depend only on η , SNR, λ , and γ_l .

Corollary .13 (Phase transitions for $\gamma_u \rightarrow 0$). In the limit $\gamma_u \rightarrow 0$, the spike direction is recovered exactly. Hence, $c(\gamma_u) \rightarrow 1$ and $D_{\mathbf{v}}(\gamma_u) \rightarrow 0$. Thus, the transition curve when $\gamma_l > 1$ simplifies to

$$(1 - \eta) \frac{1}{\gamma_l} = \eta \frac{\lambda\gamma_l(\gamma_l - 1)}{(\gamma_l - 1 + \lambda)^2} + \frac{1}{S(\gamma_l - 1)}.$$

Furthermore, when $\gamma_l < 1$, the stability boundary reduces to

$$\gamma_l^{\text{crit}} = \frac{S(1 - \eta)}{1 + S(1 - \eta)}.$$

B. Theory-experiment validation

Below, we present additional plots validating the theoretical predictions through numerical experiments. We introduce an overlap measure for the task-spike alignment θ , which can be defined through the overlap η as

$$\frac{(\mathbf{w}^{\star \top} \mathbf{v})^2}{\|\mathbf{w}^{\star}\|^2} = \eta = (\cos \theta)^2.$$

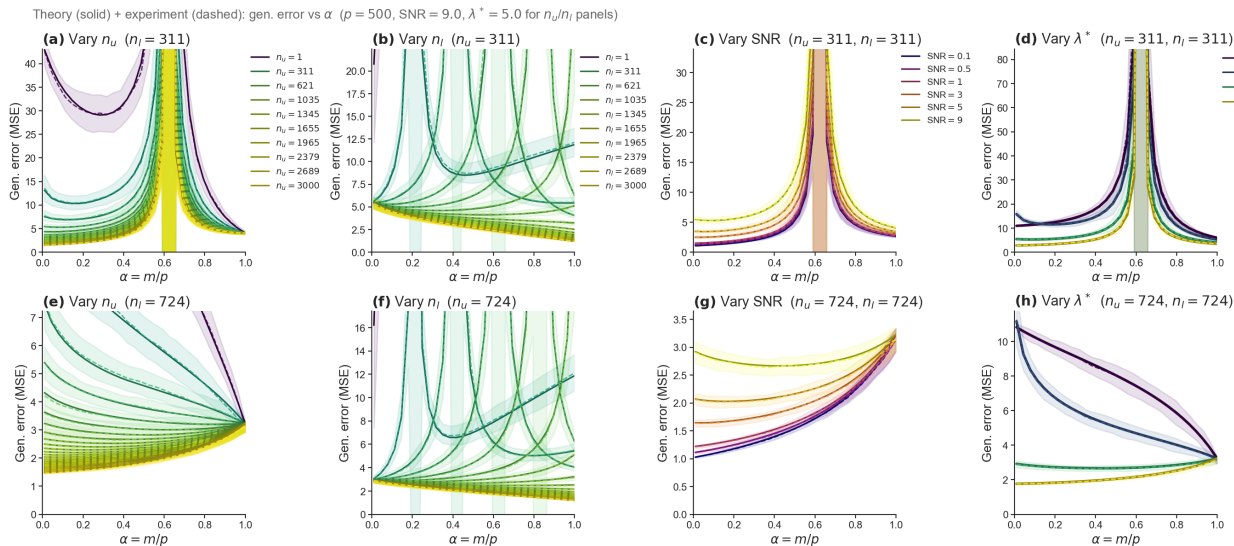


Figure 5. Generalisation error of PCR as a function of the number of retained components $\alpha = \frac{m}{p}$. Solid lines show the theoretical prediction; dashed lines show the empirical mean \pm one standard deviation across trials. Each column varies one parameter while holding the others fixed: **(a,e)** varying the pre-training dataset size n_u with n_l fixed; **(b,f)** varying the downstream dataset size n_l with n_u fixed; **(c,g)** varying the signal-to-noise ratio SNR; **(d,h)** varying the spike eigenvalue λ . The two rows correspond to two regimes: smaller (top) and larger (bottom) values of n_u and n_l . Results are shown for a rank-1 spiked covariance model with perfectly aligned signal ($\eta = 1$).

B.1. ERROR CURVES

Figure 5 illustrates the range of generalisation error curves across parameter settings. When the downstream sample size is limited ($n_l < p$), the error exhibits a double descent profile, with the peak occurring at $\alpha \approx 1/\gamma_l$. Increasing the size of the pretraining dataset n_u uniformly reduces the overall error scale (panels **(a,e)**), while varying n_l shifts the location of the interpolation threshold (panels **(b,f)**). Increasing the SNR raises the generalisation error (panels **(c,g)**), as it amplifies the contribution of signal-dependent bias. This increase is non-uniform across α : smaller values of α are affected more strongly, since a larger fraction of the signal lies outside the retained subspace. A similar non-uniform effect is observed when varying the spike strength λ (panels **(d,h)**).

We further decompose these effects by examining the estimation error components in Figure 6. The bias due to missing signal is determined by the downstream sample regime. In the overparametrized regime ($n_l < p$), this bias admits a strictly positive lower bound (panel **(a)**), reflecting the inability to fully reconstruct the signal. In contrast, when $n_l > p$, this component can vanish (panel **(e)**). Independently of this constraint, the bias decreases with α , since increasing α enlarges the subspace \mathcal{P}_m and thus increases the fraction of \mathbf{w}^* that can be represented. The leak-induced bias follows a similar decreasing trend in α , as it is controlled by $\bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}$, which shrinks as the orthogonal complement \mathcal{P}_{\perp} becomes smaller. However, its dependence on λ is non-monotonic (panels **(b,f)**): it is negligible for weak spikes (e.g. $\lambda = 1.1$), increases at intermediate values (e.g. $\lambda = 2, 5$), and decreases again for strong spikes (e.g. $\lambda = 10$). This reflects a trade-off between decreasing leak size $\bar{p}_{\perp, \mathbf{w}^*, \mathbf{v}}$ and increasing the effective component of the leak vector \mathbf{v} in the subspace $\mathbf{X}_l \mathcal{P}_m$ through $\bar{\pi}_{\mathbf{v}}$. For the parameter ranges considered, the magnitude of the leak term remains small relative to the other components. The variance behaves as expected: in the regime $n_l < p$, it decreases monotonically with λ , reflecting improved signal-to-noise separation (panels **(c,g)**).

Finally, the generalisation error (Figure 7) includes two additional contributions relative to the estimation error. The first is the additive noise variance σ^2 (panels **(c,g)**). The second is an additional bias term induced by the spike structure in the test distribution. This term combines a residual missing signal component with a correction from the leak, and it is amplified by λ . Its structure leads to a discontinuity at $\alpha = \gamma_l$ (panel **(b)**), inherited from the behaviour of the “missing signal” bias. Overall, the generalisation error closely tracks the estimation error, with the main differences appearing at small α , where these additional bias contributions are most pronounced (panels **(e,f)**).

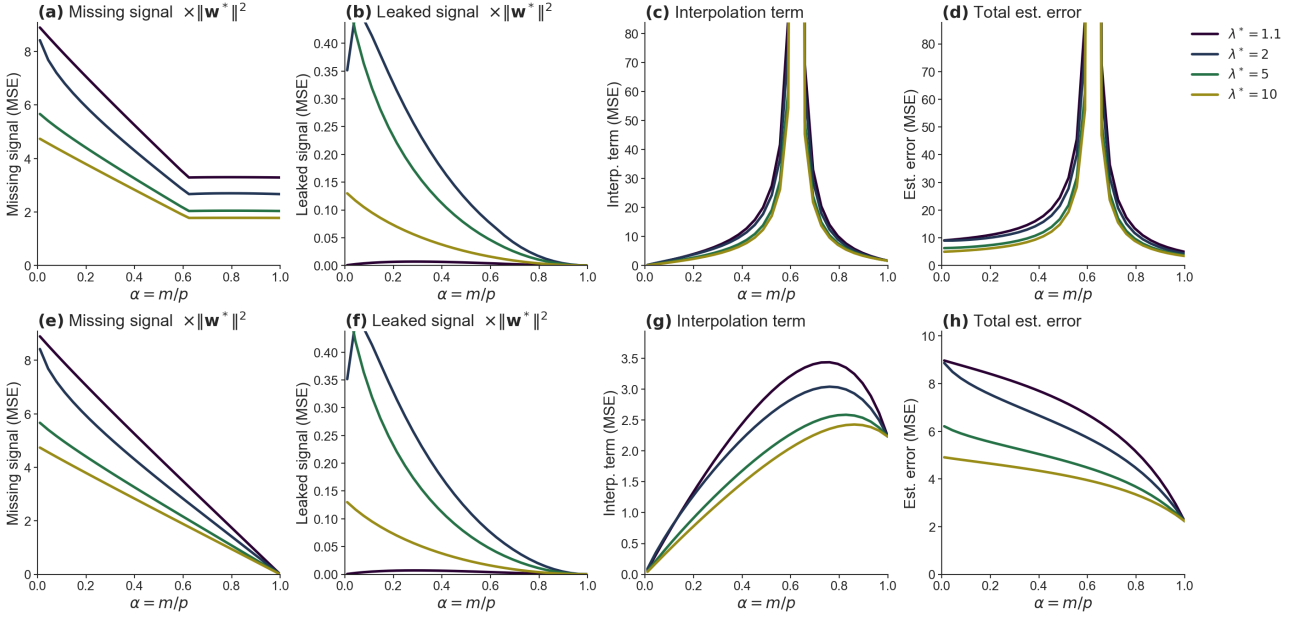


Figure 6. Components of the theoretical estimation error as a function of α for different values of λ . The total error (panels (d, h)) is decomposed into the missing signal term (panels (a, e)), the leak term (panels (b, f)), the variance term (panels (c, g)). All plots assume fixed parameters $n_u = 311$, $n_l = [311, 724]$, $\text{SNR} = \frac{\|\mathbf{w}^*\|^2}{\sigma^2}$, and a target signal misaligned with the spike eigenvector ($\eta = 0.75$). The number of features is set to $p = 500$.

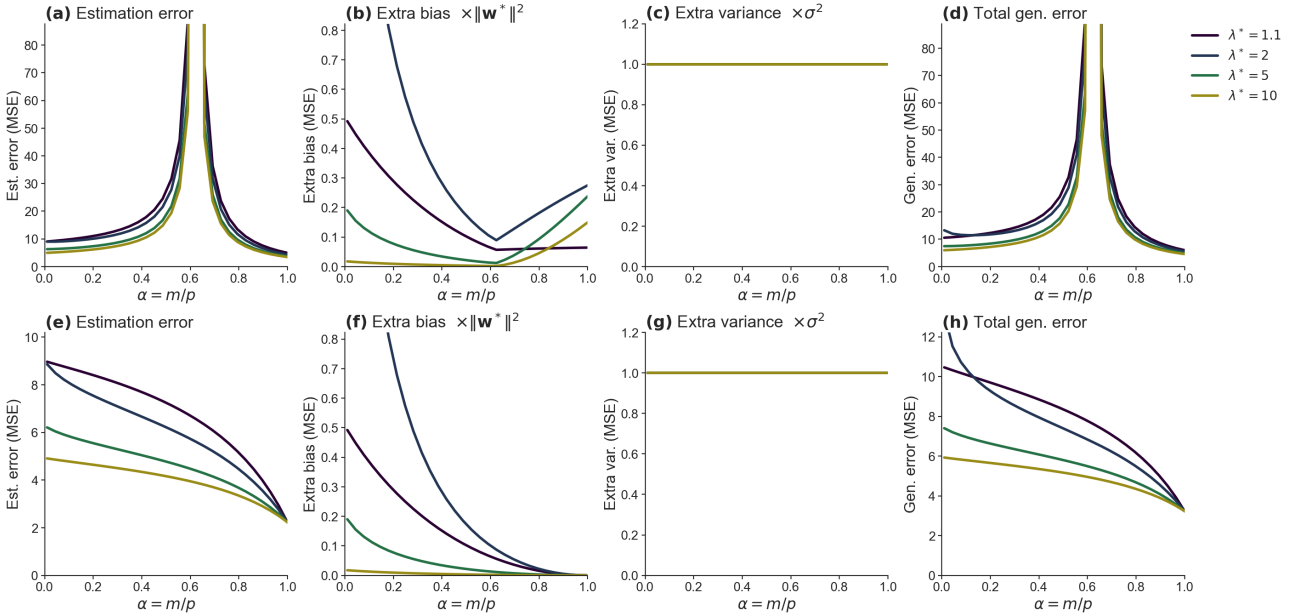


Figure 7. Components of the theoretical generalisation error as a function of α for different values of λ . The total error (panels (d, h)) is decomposed into the estimation error (panels (a, e)), the missing signal (panels (b, f)), and the leak term (c, g). All plots assume fixed parameters $n_u = 311$, $n_l = [311, 724]$, $\text{SNR} = \frac{\|\mathbf{w}^*\|^2}{\sigma^2}$, and a target signal misaligned with the spike eigenvector ($\eta = 0.75$). The number of features is set to $p = 500$.

B.2. OPTIMAL DIMENSIONALITY

In Figures 8, 9, 10 and 11, we present various phase portraits in α , reflecting their dependence on the signal-to-noise ratio (SNR), the spike strength (λ), and the task overlap. For the optimal α minimizing the training error (Figures 9 and 11),

1980 the SNR, task overlap, and spike strength have little effect. In contrast, for the optimal α minimizing the generalisation
1981 error, the phase-transition boundaries and the size of the bulk region are highly sensitive to the task overlap, SNR, and spike
1982 strength (λ). Figures 8 and 10 show that the phase-transition boundary flattens into a horizontal line at a fixed n_u (*i*) as
1983 the task-spike overlap decreases, (*ii*) as the SNR increases or (*iii*) as the spike strength decreases. At the same time, the
1984 bulk region expands along the n_u axis, thereby reducing the region of parameter space in which compression is used. This
1985 behaviour can be explained by the fact that a weaker signal is harder to detect, leading to a less accurate estimate of the prior.
1986 Consequently, exploiting this poorly-estimated prior through compression is not advantageous across many regimes, even
1987 when the amount of new data is large.

1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034

C. Additional experiments

We provide additional figures from the experiments outlined in the main text in Section 5, as well as an additional experiment directly comparing the effectiveness of pretrained regression.

C.1. COMPARISON WITH REGRESSION AND PRINCIPAL COMPONENT REGRESSION

We compare three approaches: standard regression, pretrained regression (PR) where PCA is learned on prior data and transferred to the downstream task, and principal component regression (PCR), where PCA is learned directly on the downstream data.

In the fully aligned setting ($\mathbf{w}^* \parallel \mathbf{v}$, $\cos \theta = 1$), Figure 12 (left) shows that incorporating a PCA step consistently improves performance over standard regression. Moreover, when the pretraining dataset is larger than the downstream dataset ($n_u > n_l$), PR outperforms PCR, reflecting the more accurate estimation of the signal subspace from additional data.

In the misaligned setting ($\cos \theta = 1/\sqrt{2}$, right panel in Figure 12), the advantage of PR is restricted to regimes with large n_u and small n_l . As n_l increases, PCR becomes comparable to PR, and both methods approach the performance of standard regression (within a small error gap). This reflects the diminishing benefit of transferring a pretrained subspace when sufficient downstream data is available to estimate the relevant structure directly.

C.2. ADDITIONAL AUTOENCODER EXPERIMENTS

Figure 13 extends panels (a-d) of Figure 4. The top row shows simulations that directly match our theoretical setting (PCA followed by regression), while the middle and bottom rows replace PCA with unsupervised representation learning using a linear and a nonlinear one-hidden-layer autoencoder, respectively. The linear autoencoder closely reproduces the theoretical predictions in both the spiked identity and general covariance settings, confirming that it effectively recovers the same subspace as PCA in these regimes.

We next examine the effect of reducing the SNR in Figure 14. The phase boundaries shift in the same qualitative manner as predicted by the theory, where the optimal α is low across a larger part of the (n_u, n_l) space, indicating that the dependence on SNR is captured beyond the exact PCA setting.

Finally, we analyse deviations from PCA in low-sample regimes in Figure 15. We compare the theoretical overlap $\mathcal{P}_m \mathbf{v}$ with the empirical overlap obtained by passing \mathbf{v} through trained linear and nonlinear autoencoders with varying bottleneck dimension. The linear autoencoder deviates from PCA only when $n_u < p$, while the nonlinear autoencoder exhibits larger discrepancies across a broader range of n_u .

Despite these deviations, the optimal bottleneck dimension that minimizes both training and generalisation error remains close to $\alpha \approx 1$ when $n_u < p$. As a result, the performance of linear autoencoder-based pipelines is still well predicted by the theory across all n_u . In contrast, the nonlinear autoencoder exhibits a persistent mismatch for $n_u > p$, which only diminishes for sufficiently large pretraining datasets (empirically around $n_u \approx 3p$). The magnitude of this mismatch also depends on SNR, as seen by comparing Figures 13 and 14.

C.3. ADDITIONAL TRANSFORMER EXPERIMENTS

We first consider the setting where PCA is applied to representations extracted from the downstream task (Figure 16), rather than from pretraining data as in Figure 4(e). In this regime, compression remains beneficial for generalisation when the downstream dataset is small, with an optimal intermediate dimension m . The dependence on the pretraining dataset size n_u is weaker than in the pretraining-PCA setting, but still observable. This likely reflects the evolution of the representation spectrum during training, with increasing spike size (effective λ) as shown in Figure 18.

We next return to applying PCA to pretraining data and report validation accuracy across a broader range of downstream dataset sizes (Figure 17). For small downstream datasets, an intermediate compression level is optimal, particularly for models trained for longer durations. In contrast, for large downstream datasets, performance is maximized without compression ($m \approx p$). At the smallest dataset sizes, the optimal m is less stable across runs.

Neither experimental setup provides a direct match to the theoretical model. In particular, the pretraining and downstream datasets differ not only in semantic structure but also in sequence length, which affects the learned representations and limits the transferability of PCA computed on pretraining data. Despite this mismatch, both PCA-based approaches match or

outperform the no-compression baseline (Figure 19).

D. Implementation details

D.1. FIGURE 1

Panel (a) Diagram of the model made with `figma`.

Model. We define the signal-to-noise ratio as $\text{SNR} = \frac{\|\mathbf{w}^*\|^2}{\sigma^2}$ and the alignment between the regression target and the spike as $\eta = \cos \theta$. Throughout, we fix $p = 500$ and $\sigma^2 = 1$.

Panels (b-c) (generalisation and training error vs α). We fix $n_u \approx n_l \approx 300$, $\text{SNR} = 9.0$, and $\eta = 1.0$. The parameter α is swept over 30 values in $[0.01, 1.0]$ for each of four spike strengths $\lambda \in \{1.1, 2.0, 5.0, 10.0\}$. Theoretical predictions (solid curves) are compared against empirical averages over multiple realisations (dashed curves, shading corresponds to 1 standard deviation).

D.2. FIGURE 2

Model. We define the signal-to-noise ratio as $\text{SNR} = \frac{\|\mathbf{w}^*\|^2}{\sigma^2}$ and the alignment between the regression target and the spike as $\eta = \cos \theta$. Throughout, we fix $p = 500$ and $\sigma^2 = 1$.

Panel (a) (minimum generalisation error). For $\text{SNR} = 9.0$, $\lambda^* = 5.0$, and $\eta = 1.0$, we plot $\min_{\alpha} E_{\text{gen}}(\alpha)$ over a 100×100 grid with $n_u, n_l \in [1, 3000]$. Dashed white lines mark $n_u = p$ and $n_l = p$ ($p = 500$). The colour scale is logarithmic.

Panel (b) (gain from tuning α relative to $\alpha = 1$). Under the same setting, we compute $E_{\text{gen}}(\alpha = 1) - E_{\text{gen}}(\alpha^*)$ over the same grid. Positive values indicate regions where using the full prior representation is suboptimal and tuning α reduces generalisation error. The colour scale is logarithmic.

Panel (c) (gain from tuning α relative to α_{\min}). Under the same setting, we compute $E_{\text{gen}}(\alpha = \alpha_{\min}) - E_{\text{gen}}(\alpha^*)$, where $\alpha_{\min} \approx 0.01$ is the smallest value in the sweep ($m = 1$ principal component retained). This quantifies the benefit of tuning α upward from the minimal representation. The colour scale is logarithmic.

Panel (d) (marginal substitution rate). Under the same setting, we compute the ratio $(\partial E_{\text{gen}} / \partial n_u) / (\partial E_{\text{gen}} / \partial n_l)$ evaluated at the optimal α^* , with gradients approximated numerically via central differences on the 100×100 grid. Values above 1 (red) indicate that collecting an additional unlabelled sample n_u reduces generalisation error more than collecting a labelled sample n_l ; values below 1 (blue) indicate the converse. The colour scale is centred at 1.

Panel (e) (generalisation-optimal α^*). Under the same setting, we plot the value of $\alpha^* = \arg \min_{\alpha} E_{\text{gen}}(\alpha)$ over the same grid, normalised so that $\alpha \in [0, 1]$.

D.3. FIGURE 3

Model. We define the signal-to-noise ratio as $\text{SNR} = \frac{\|\mathbf{w}^*\|^2}{\sigma^2}$ and the alignment between the regression target and the spike as $\eta = \cos \theta^2$. Throughout, we fix $p = 500$ and $\sigma^2 = 1$.

Panel (a) (optimal α^* heatmap). For $\text{SNR} = 9.0$, $\lambda = 5.0$, and $\eta = 1.0$, we compute the theoretical optimal representation size α^* that minimises generalisation error over a 100×100 grid, where n_u and n_l each vary from 1 to 3000. Analytical phase-transition boundaries (smooth transition and sharp transition) are overlaid.

Panel (b) (phase portrait SNR). At fixed $\lambda = 5.0$ and $\eta = 1.0$, we plot the smooth and sharp transition boundaries in the (n_u, n_l) plane for each of the six SNR values $\{0.1, 0.5, 1.0, 3.0, 5.0, 9.0\}$. Increasing SNR progressively reduces the region in which full-rank representations ($\alpha \approx 1$) are suboptimal.

Panel (c) (phase portrait spike strength). At fixed SNR = 9.0 and $\eta = 1.0$, we show the corresponding transition curves for $\lambda \in \{2, 3, 5, 7, 10\}$. Stronger spikes shift the phase boundaries.

Panel (d) (phase portrait alignment). At fixed SNR = 9.0 and $\lambda = 5.0$, we plot the transition curves for $\eta \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$. All panels are evaluated over $n_u, n_l \in [1, 3000]$.

D.4. FIGURE 4

Autoencoders

Data generation. Both covariances use $p = 200$ features, noise standard deviation $\sigma = 1$, SNR = 9, $\|\mathbf{w}^*\| = 3$.

- *Spiked identity* ($\lambda = 5.0$): $\Sigma = \mathbf{I}_p + (\lambda - 1)\mathbf{v}\mathbf{v}^\top$; true weight vector aligned with the spike $\mathbf{w}^* = \|\mathbf{w}^*\|\mathbf{v}$, where \mathbf{v} is the top eigenvector of Σ .
- *Spiked Toeplitz* ($\lambda = 5.0$): Σ is a Toeplitz matrix with $\rho = 0.5$, with λ added to the top eigenvalue; \mathbf{w}^* again aligned with the top eigenvector.

Each grid point (n_u, n_l) averages five independent replications, with $n_u, n_l \in [1, 1200]$ each swept over 20 uniform values.

Models. For each (n_u, n_l) pair, $\alpha \in \{0.01, 0.05, \dots, 1.0\}$ (20 values) sets the latent dimension $m = \lfloor \alpha p \rfloor$. The three models below are evaluated.

1. *PCA + Regression (PR)*: Eigenvectors of $\mathbf{S}_u = \mathbf{X}_u^\top \mathbf{X}_u / n_u$ are computed; \mathbf{X}_l is projected onto the top- m eigenvectors and OLS is applied in the latent space.
2. *Linear AE*: Encoder ($\mathbb{R}^p \rightarrow \mathbb{R}^m$) and linear decoder ($\mathbb{R}^m \rightarrow \mathbb{R}^p$) trained by minimizing MSE on \mathbf{X}_u ; the frozen encoder maps \mathbf{X}_l to features for OLS.
3. *Nonlinear AE*: Same as the linear AE but with a sigmoid nonlinearity applied at the encoder, $\mathbf{z} = \sigma(\mathbf{W}_{\text{enc}} \mathbf{x})$.

All autoencoders are implemented in JAX/Equinox and trained with Adam (lr = 10^{-3} , 20,000 epochs, batch size 500). Weights are initialized as $\mathbf{W}_{i,j} \sim \mathcal{N}(0, 1/\sqrt{d_{\text{in}}})$.

Panel descriptions. Each panel is a 20×20 heatmap with n_u on the x -axis, n_l on the y -axis, and color encoding $\alpha^* \in [0, 1]$ (the value minimizing mean loss over the five replications).

The autoencoder component of the figure is organized as 2 rows \times 2 columns:

	Linear AE	Nonlinear AE
Spiked identity	(a)	(b)
Spiked Toeplitz	(c)	(d)

Columns compare linear vs nonlinear AEs; rows contrast spiked identity and spiked Toeplitz covariance structures; each panel shows generalisation-optimal α^* . A single shared colorbar applies to all panels.

LLMs

Model and representations. All experiments use Pythia-70M-deduped (Biderman et al., 2023) (70M parameters, trained on the deduplicated Pile (Gao et al., 2020; Biderman et al., 2022)). We evaluate 10 checkpoints at steps 2, 8, 32, 128, 512, 2000, 5000, 20000, 80000, and 143,000. Representations are last-token hidden states from the final layer of 64-token truncated inputs, yielding vectors in \mathbb{R}^{512} .

Downstream task. Binary sentiment classification on SST-2 from the GLUE benchmark (Wang et al., 2018). Fixed 872-sentence validation set; training sizes $n_l \in \{100, 250, 500, 1000, 2000, 4000, 10000, 20000, 40000, 67349\}$, n_l training examples are randomly chosen out of the maximum 67349 for each seed.

PCA. We sweep $m \in \{2, 3, 4, 7, 11, 16, 25, 39, 60, 92, 142, 218, 334, 512\}$ (approximate \log_2 grid), and consider the two conditions below.

1. *Downstream PCA* (Figure 16): basis fit on the n_l labelled training representations;
2. *Pretrain-PCA* (Figure 4(e-g)): basis fit on up to 51,200 Pile sequences (50 batches \times 1,024 sequences) from the same checkpoint, using `torch.linalg.eigh`. For earlier checkpoints, the basis is fit on $1024 \times n_{\text{checkpoint}}$ sequences.

Linear probe. Projected representations are fed to a 2-output linear layer trained with cross-entropy loss using Adam ($\text{lr} = 10^{-3}$), for up to 10,000 epochs with early stopping. Linear probes were randomly initialized with small weights for each seed.

Seeds and aggregation. Each (checkpoint, n_l , m) configuration is run with 3 random seeds (affecting probe initialization and downstream data shuffling). Figures show mean \pm std; optimal- m heatmaps report the mean per-seed argmax.

Panel descriptions.

- **(e)** Heatmap for optimal m under pretrain-PCA, selected by validation accuracy.
- **(f)** Eigenvalue histogram at step 143,000, computed from the 872-sentence SST-2 validation covariance; log-density on the y -axis. The full set of 5 checkpoints is shown in Figure 18.
- **(g)** Validation accuracy vs m (x -axis in \log_2 scale) under pretrain-PCA for $n_l = 2000$, with curves colored by pretraining progress.

D.5. COMPUTE

Theory curves. Theory curves were computed on a single CPU (Intel Xeon E5-2680 v4 @ 2.40 GHz, no GPU required). Each configuration sweeps a 100×100 grid of (n_u, n_l) values over 30 values of α , amounting to approximately 300,000 evaluations of closed-form expressions. The full sweep covers 24 configurations, $\lambda^* \in \{1.1, 2.0, 5.0, 10.0\} \times \text{SNR} \in \{0.1, 0.5, 1.0, 3.0, 5.0, 9.0\}$, with alignment $\eta = 1.0$, totaling approximately 7.2 million evaluations. When executed in parallel, the full sweep completes in approximately 2h. Peak RAM usage was approximately 2 GB per configuration.


Simulations. Simulations were distributed over a SLURM cluster. Each job was allocated 4 CPUs and 30 GB of RAM. The full grid $6 \text{ SNR} \times 4 \lambda_1 \times 4 \lambda_2 \times 30 n_b$ produced 2,880 SLURM jobs, each processing 30 values of n_l with 100 independent replications at $p = 500$. Each replication involves an eigendecomposition and pseudoinverse of a $p \times p$ matrix for each of the 30 values of α . When executed in parallel, the full sweep completes in approximately 10h.

Autoencoder experiment. Simulations were distributed over a SLURM cluster. Each job was allocated 1 GPU, 8 CPUs, and 16 GB of RAM. The full grid $20 n_{\text{base}} \times 20 n_{\text{task}}$ produced 400 SLURM jobs, each processing 20 values of α with 5 independent replications at $p = 200$. Each replication trains a linear and a nonlinear autoencoder for 20,000 epochs for each value of α , followed by OLS regression in the learned latent space. The individual jobs completed in 10-50min depending on parameters.

LLM probe experiment. Simulations were distributed over a SLURM cluster. Each job was allocated 1 GPU, 16 CPUs, and 32 GB of RAM. Two sets of 10 SLURM jobs were submitted: one per checkpoint of Pythia-70M-deduped for the downstream-PCA condition and one for the pretrain-PCA condition. Each job sweeps over 10 training-set sizes and 3 seeds, fitting a linear probe for up to 10,000 epochs with early stopping across up to 14 PCA dimensionalities m .


D.6. LLM USAGE

Large language models (LLMs) were used throughout this research project to assist with writing, editing, and formatting the manuscript, including improving clarity and refining \LaTeX expressions. LLMs were also used as an aid for understanding and discussing technical concepts, as well as for debugging and drafting portions of the experimental and simulation code. All scientific results, analyses, and conclusions were verified and validated by the authors.



2303 *Figure 8.* Optimal value of α that minimizes the generalisation error, shown as a function of SNR and spike alignment η . We fix $\lambda = 2$.
2304 Red lines indicate the approximate phase transitions.

2305
2306
2307
2308
2309



2358 *Figure 9.* Optimal value of α that minimizes the training error, shown as a function of SNR and spike alignment η . We fix $\lambda = 2$.

2359
2360
2361
2362
2363
2364

2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419

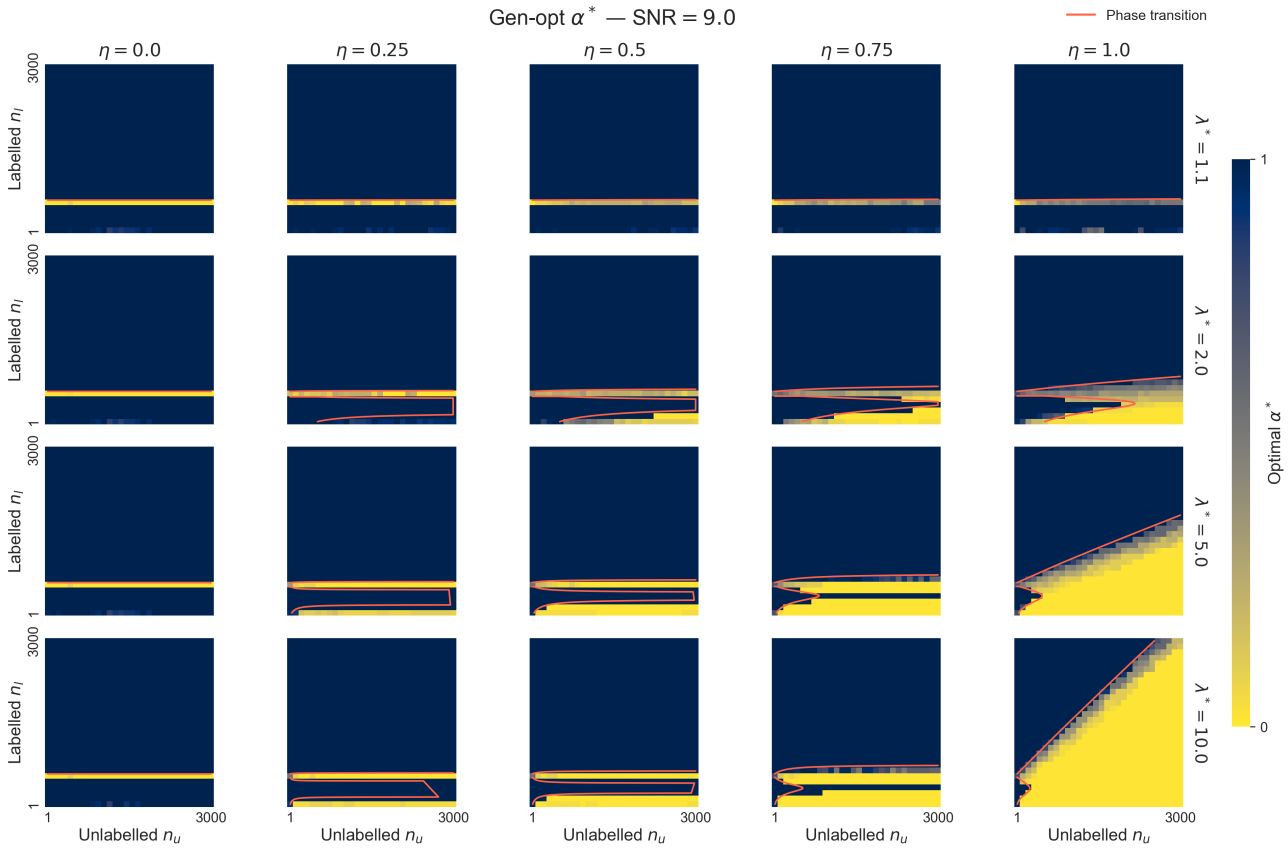


Figure 10. Optimal value of α that minimizes the generalisation error, shown as a function of spike strength λ and spike alignment θ with \mathbf{w}^* . We fix SNR = 9. Red lines indicate the approximate phase transitions.

2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474




Figure 11. Optimal value of α that minimizes the training error, shown as a function of spike strength λ and spike alignment θ with \mathbf{w}^* . We fix $\text{SNR} = 9$.

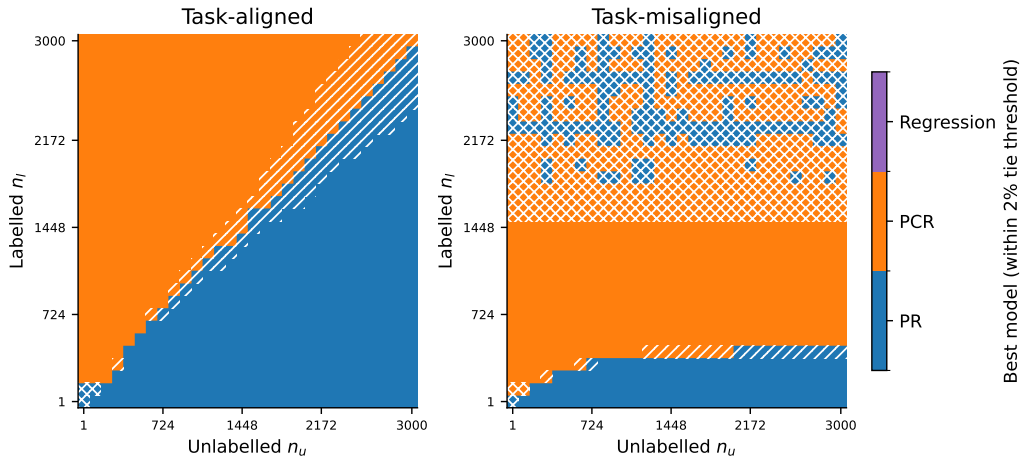


Figure 12. Heatmaps over the (n_u, n_l) grid ($p = 500$, $\text{SNR} = 1.8$, $\lambda = 5$) showing which method – standard regression (violet), pretrained regression (PR, blue) or principal component regression (PCR, orange) – achieves lower generalisation error. Forward hatching ($//$) marks regions where PR and PCR are tied (within 2% error); cross hatching (\times) marks regions where all three methods (PR, PCR, and standard regression) are tied. When the task is aligned with the pretraining signal, PR wins for $n_u > n_l$ and PCR wins otherwise, with the boundary tracking the diagonal. Under misalignment, PCR dominates across most of the grid; PR retains an advantage when $n_u \gg n_l$.

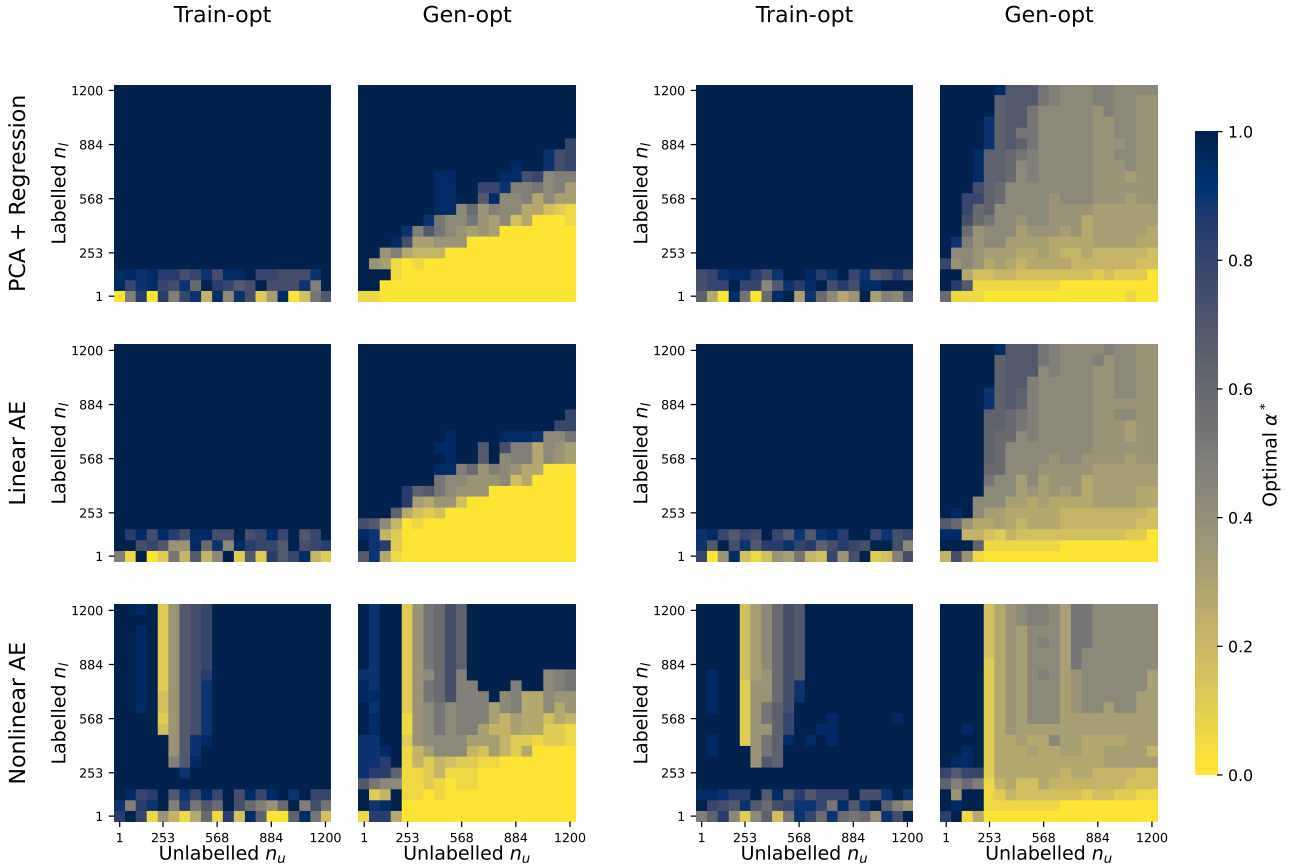


Figure 13. Extended version of panels (a-d) in Figure 4: optimal bottleneck size in autoencoders on inputs with different covariance structure. We compare PCA on pretraining data (top row) corresponding to our theoretical setup, a linear autoencoder (middle row) and a non-linear autoencoder (bottom row) trained on Gaussian data, with varying bottleneck size m ($\alpha = m/p$). The input covariance on the first two columns is $\Sigma = \mathbf{I} + \lambda \mathbf{v}\mathbf{v}^\top$, while in the last two columns the covariance is a Toeplitz matrix \mathbf{H} with a spike, i.e., $\Sigma = \mathbf{H} + \lambda \mathbf{v}\mathbf{v}^\top$, where $H_{i,j} = \rho^{|i-j|}$ with $\rho = 0.5$. The autoencoder results match those of PCA, and the identity vs general covariance exhibit a similar phenomenology. We fix $\text{SNR} = 9.0$, $\lambda = 5.0$.

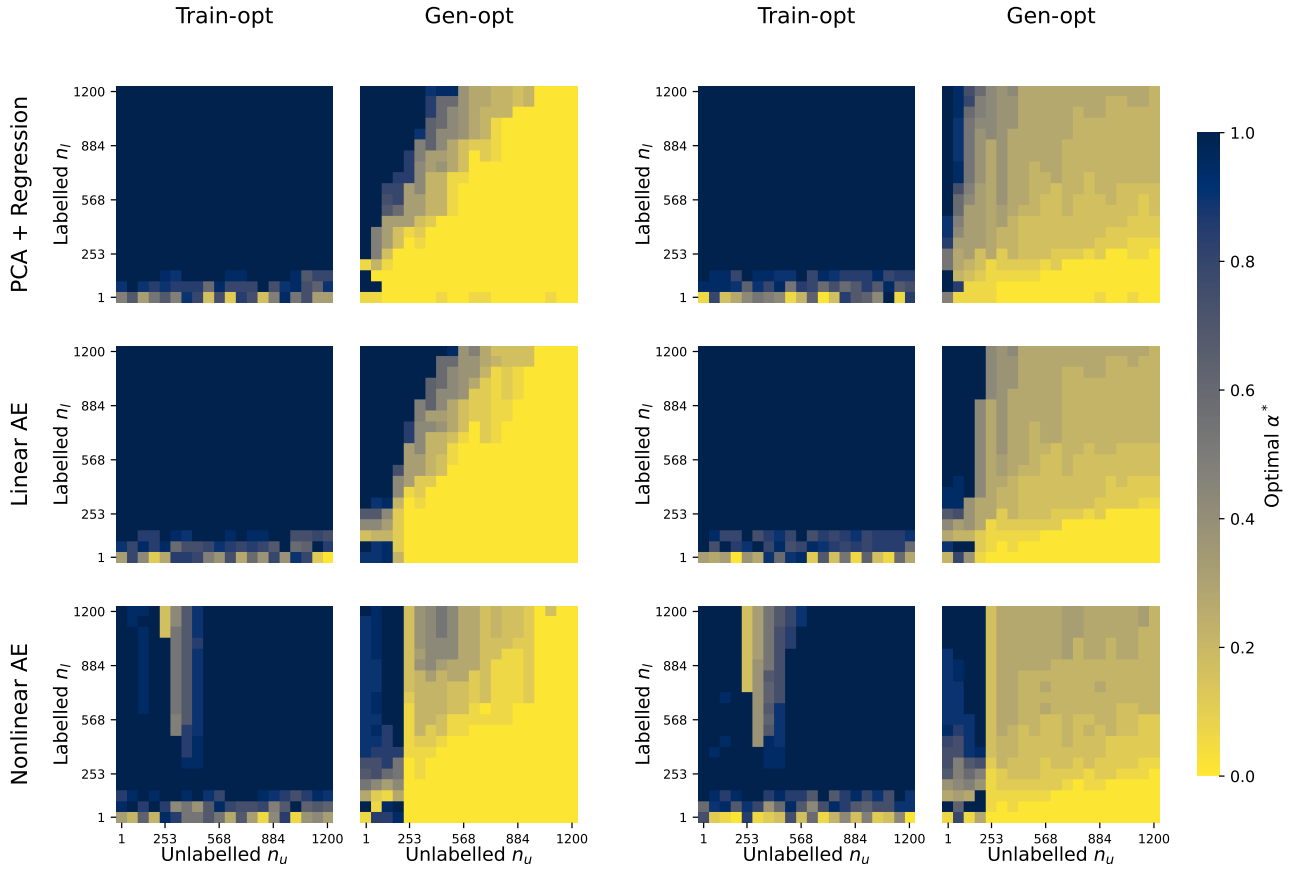


Figure 14. Effect of SNR on optimal bottleneck size in autoencoders on inputs with different covariance structure. We compare PCA on pretraining data (top row) corresponding to our theoretical setup, a linear autoencoder (middle row) and a non-linear autoencoder (bottom row) trained on Gaussian data, with varying bottleneck size m ($\alpha = m/p$). The input covariance on the first two columns is $\Sigma = \mathbf{I} + \lambda \mathbf{v}\mathbf{v}^\top$, while in the last two columns the covariance is a Toeplitz matrix \mathbf{H} with a spike, i.e., $\Sigma = \mathbf{H} + \lambda \mathbf{v}\mathbf{v}^\top$, where $H_{i,j} = \rho^{|i-j|}$ with $\rho = 0.5$. The autoencoder results match those of PCA, and the identity vs general covariance exhibit a similar phenomenology. We fix $\text{SNR} = 1.8$, $\lambda = 5.0$.




Figure 15. Spike direction recovery as a function of sample ratio n_u/p for linear and nonlinear autoencoders. Each panel shows the overlap between the singular vectors of the encoder weights and the population spike \mathbf{v} as a function of the sample ratio n_u/p , for a fixed latent dimension $m \in \{5, 25, 50, 75, 100\}$. Data are generated from a spiked covariance model in $p = 100$ dimensions with spike strength $\lambda = 5$. Solid lines show empirical means over 5 random seeds; shaded bands show ± 1 standard deviation. The dotted line shows the theoretical PCA overlap $\bar{p}_{\mathbf{v}, \mathbf{v}}$ derived from Lemma 7. Both linear and nonlinear autoencoders show a mismatch in low sample regimes, with a more evident mismatch in nonlinear autoencoders and in autoencoders with larger bottlenecks. For large sample regimes, linear and nonlinear autoencoders match PCA well, and show the phase transitions predicted by our theory (see panels (a-d) in Figure 4 discussed in Section 5.1, and Figures 13-14 discussed in Appendix C.2).

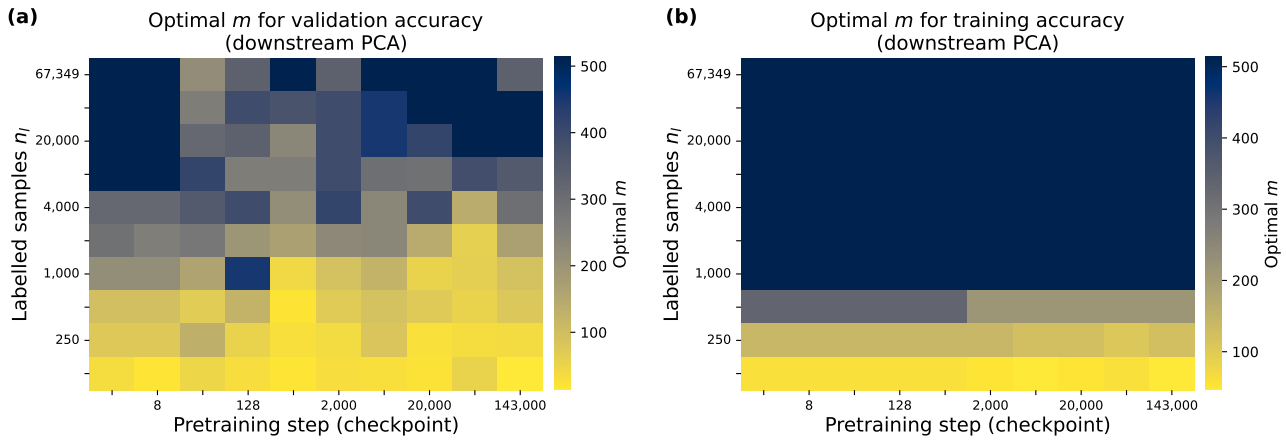


Figure 16. Optimal m using PCA on representations of the downstream task. We use the same layout as Figure 4(e), but the PCA basis is fit on the downstream labelled set rather than on the pretraining corpus. (a) Optimal m for validation accuracy; (b) optimal m for training accuracy. Compared to doing PCA on the pretraining corpus (Figure 4(a)), downstream PCA tends to favor larger m and the dependence on the checkpoint is weaker.

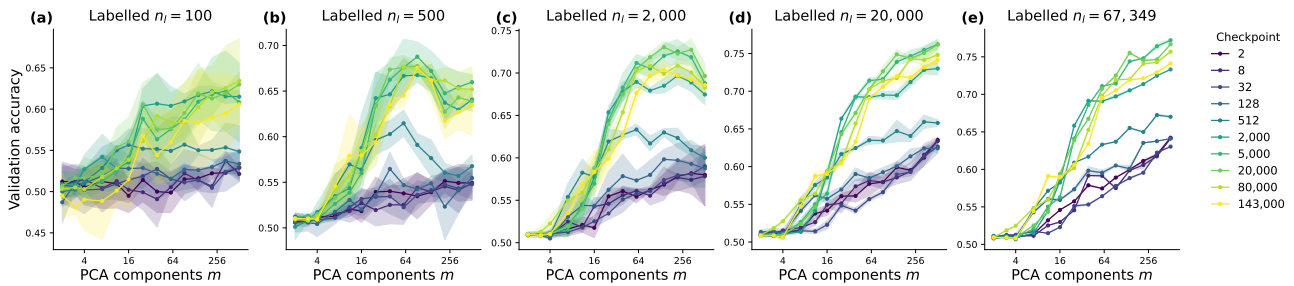


Figure 17. Validation accuracy vs number of PCA components m across five downstream task sizes. Note that the values of m are on a log scale. Each panel shows a different value of n_l spanning the full range from $n_l = 100$ to $n_l = 67,349$; curves are coloured by pretraining checkpoint. Shaded bands show the standard deviation over three random seeds. For smaller tasks, the curves have a clear peak and become progressively more monotone as n_l grows, reflecting that small tasks benefit from aggressive compression while large tasks can exploit finer-grained directions in representation space.

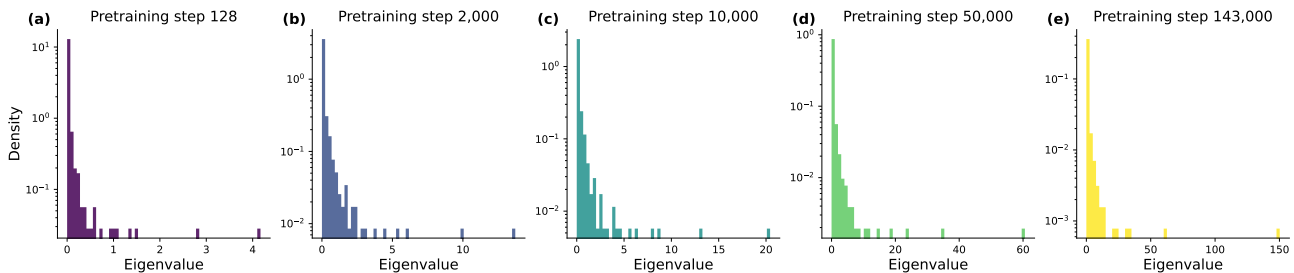


Figure 18. Eigenvalue spectra of Pythia-70M-deduped last-token last-hidden-layer representations at five pretraining checkpoints. Each panel shows the empirical distribution of covariance eigenvalues computed from the representation of the `sst2` task. Throughout pretraining, the spectrum shows a clear bulk and a few outlying eigenvalues. As pretraining progresses (steps 128 to 143,000), the outlying eigenvalues become more and more prominent.

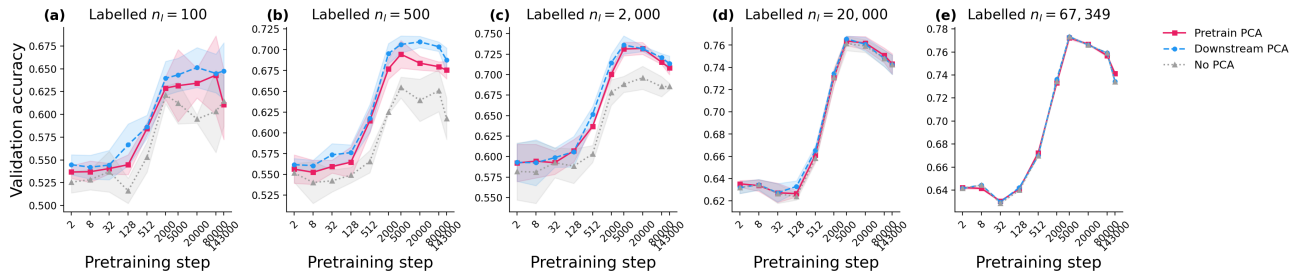


Figure 19. Validation accuracy vs pretraining checkpoint for three probing conditions across five downstream task sizes. Each panel shows a different downstream task size n_l ; the x -axis is the pretraining step on a log scale. Three curves are compared: Pretrain PCA (pink, solid) fits the PCA basis on `File` representations from the same checkpoint and selects the best m ; Downstream PCA (blue, dashed) fits PCA on the downstream task representations (`st2`) and selects the best m ; No PCA (grey, dotted) uses the full-dimensional representations (largest available m). For each curve, the displayed value is the mean over three seeds at the per-seed optimal m , with the shaded area corresponding to 1 standard deviation. The gap between PCA and the no-PCA baseline is largest at small and intermediate n_l , where dimensionality reduction acts as regularization. Pretrain PCA matches downstream PCA for intermediate and large n_l , with the mismatch coming from the difference in the type of data available in the pretraining and downstream datasets.