PCA++: How Uniformity Induces Robustness to Background Noise in Contrastive Learning

Mingqi Wu
McGill University
Mila
mingqi.wu@mail.mcgill.ca

Qiang Sun
University of Toronto
MBZUAI
qsunstats@gmail.com
Corresponding author

Archer Y. Yang
McGill University
Mila
archer.yang@mcgill.ca
Corresponding author

Abstract

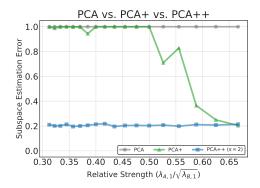
High-dimensional data often conceal low-dimensional signals beneath structured background noise, limiting standard PCA. Motivated by contrastive learning, we address the problem of recovering shared signal subspaces from positive pairs-paired observations sharing the same signal but differing in background. Our baseline, PCA+, uses alignment-only contrastive learning and succeeds when background variation is mild, but fails under strong noise or high-dimensional regimes. To address this, we introduce PCA++, a hard uniformity-constrained contrastive PCA that enforces identity covariance on projected features. PCA++ has a closed-form solution via a generalized eigenproblem, remains stable in high dimensions, and provably regularizes against background interference. We provide exact highdimensional asymptotics in both fixed-aspect-ratio and growing-spike regimes, showing uniformity's role in robust signal recovery. Empirically, PCA++ outperforms standard PCA and alignment-only PCA+ on simulations, corrupted-MNIST, and single-cell transcriptomics, reliably recovering condition-invariant structure. More broadly, we clarify uniformity's role in contrastive learning—showing that explicit feature dispersion defends against structured noise and enhances robustness.

1 Introduction

Real-world data often hides a simple, low-dimensional signal beneath layers of structured noise and random variation. In genomics, batch effects blur true biological differences [32, 21]. In medical imaging, scanner differences obscure true clinical signals [18]. In finance, the performance of individual assets is masked by market-wide trends [17, 9]. Yet standard Principal Component Analysis (PCA) indiscriminately captures all dominant directions, failing to distinguish the signal of interest from unwanted backgrounds.

Alignment-only contrastive learning and PCA+. Contrastive learning [20, 37, 11, 13] offers a natural remedy. By comparing paired datasets $X, X^+ \in \mathbb{R}^{n \times d}$ —each with n samples in d dimensions that share the same signal but experience different backgrounds—we can cancel out unwanted variation. Motivated by this idea, we introduce PCA+, a vanilla contrastive PCA method that forms a "contrastive" covariance $S_n^+ = \frac{1}{2n}(X^TX^+ + X^{+T}X)$ from positive pairs, and applies ordinary PCA to S_n^+ . At the population level, PCA+ perfectly recovers the shared signal subspace. In finite-sample, low-dimensional regimes, PCA+ again succeeds. However, as background spikes (eigenvalue $\sqrt{\lambda_{B,1}}$) increase relative to signal spikes (Figure 1, left) or the ambient dimension d grows relative to n (Figure 1, right), its estimated directions can drift into background-dominated components and miss the true signal.

39th Conference on Neural Information Processing Systems (NeurIPS 2025).



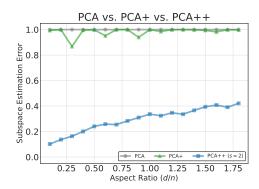


Figure 1: Subspace estimation error for standard PCA, PCA+, PCA++. Results are for Example 3.3. **Left:** (varying relative strength of the signal $\lambda_{A,1}/\sqrt{\lambda_{B,1}}$) As background strength grows, PCA+ deteriorates sharply while PCA++ keeps its error uniformly low. **Right:** (varying aspect ratio d/n) Across all regimes, PCA++ outperforms both PCA and PCA+.

Hard uniformity constraint and PCA++. To guard against this failure mode, we propose PCA++, which augments the contrastive objective with a "hard uniformity" constraint [50]: the projected features must have identity covariance. Equivalently, PCA++ solves

$$\max_{V \in \mathbb{R}^{d \times k}} \operatorname{tr}(V^{\top} S_n^+ V), \qquad \text{s.t.} \qquad V^{\top} S_n V = I_k, \tag{1.1}$$

where $S_n = \frac{1}{n}X^TX$ is the usual sample covariance. This constraint—forcing the learned features to be uniformly distributed—acts like an "even-spread" regularizer, protecting the signal representation from distortion by strong background noise. We show that PCA++ admits a closed-form solution via a *generalized eigenproblem* and remains stable (with a low-rank truncation of S_n) even when $d \gg n$. Empirically, PCA++ maintains tight alignment with the ground-truth signal across a wide range of d/n, outperforming both standard PCA and PCA+ (Figure 1).

Theoretical guarantees. We analyze PCA+ and PCA++ under a tractable contrastive factor model where each paired observation decomposes as orthogonal signal and background. Our main results include: (1). Population—level consistency of PCA+, plus finite—sample bounds on its subspace error—and a precise characterization of when strong background spikes overwhelm it. (2). Exact high-dimensional asymptotics for PCA++ in two regimes: (i) Fixed aspect ratio $d/n \rightarrow c > 0$, under a Baik-Ben Arous-Péché (BBP) detectability condition [5], we derive a closed-form limit for the subspace error; (ii) Growing-spike regime: when each spike scales with d/n, the limiting subspace estimation error simplifies to a function of the weakest signal's effective signal-to-noise-ratio. These analyses quantify how the uniformity constraint regularizes against background interference.

Broader impact on contrastive learning. Our work also sheds light on an open question in contrastive learning theory: the practical role of uniformity. Recent theoretical results [50, 39] show that contrastive losses naturally encourage uniformity while preserving meaningful class structures. Yet, the precise theoretical reason uniformity improves downstream performance–particularly in noisy, high-dimensional environments–remained unclear. We rigorously show how an explicit uniformity constraint can defend against structured noise in high-dimensional regimes–clarifying why uniformity enhances robustness and generalization.

In summary, our contributions to the community include:

- A novel contrastive PCA framework that leverages paired observations to perfectly isolate shared signal subspaces at the population level and provably bounds finite-sample subspace recovery error under a tractable linear factor model.
- 2. A hard uniformity–constrained PCA++ algorithm that enforces identity covariance on projected features, admits a closed-form generalized eigenproblem solution, and remains stable even when background eigenvalues and ambient dimension grow arbitrarily large.
- 3. Exact high-dimensional asymptotic characterizations of PCA++ in both fixed aspect-ratio and growing-spike regimes, yielding closed-form limits for subspace estimation error that quantify how uniformity regularizes against structured noise.

4. A theoretical explanation for uniformity's power in contrastive learning, rigorously showing how enforcing feature dispersion defends signal recovery from strong background interference and clarifies uniformity's role as a robust regularizer.

1.1 Related work

Foundations of contrastive learning. Early theory framed InfoNCE through mutual information (MI) maximization [37, 3], and initial analyses linked the contrastive objective to downstream class separability guarantees [43]. But MI alone proved insufficient—tighter MI bounds sometimes hurt performance [34, 48]. A major advance came with the *alignment—uniformity* framework of Wang and Isola [50], showing that InfoNCE naturally pulls positive pairs together while spreading features uniformly on the hypersphere. This insight inspired methods like Barlow Twins [52] and VICReg [7], which explicitly shape feature covariances, and it was further extended by Chen et al. [12], who parameterized alignment and uniformity terms directly. More recently, Parulekar et al. [39] proved that InfoNCE promotes uniformity in the population limit, and other work has highlighted how inductive biases steer the final representations [42, 22]. Yet a precise, theory-grounded explanation for why uniformity bolsters robustness—especially under structured noise—remains open, and it motivates our current study.

Spectral views. Another strand casts contrastive learning as a form of spectral learning. The spectral contrastive loss [23, 24] directly manipulates feature covariances, and in fact it and InfoNCE share the same population solution [29]. This spectral view connects CL to graph-based dimensionality reduction (e.g. t-SNE and spectral clustering) [47].

Linear models and high-dimensional analysis. In parallel, tractable linear and asymptotic models have shed light on CL's strengths. For instance, Ji et al. [28] showed that CL recovers latent signal directions more reliably than classic unsupervised methods, while Bansal et al. [6] used Gaussian mixtures to prove InfoNCE can identify optimal low-dimensional subspaces. Our work also builds on a tractable model, embedding a linear contrastive factor model in high-dimensional random matrix theory to isolate how a hard uniformity constraint protects signal recovery under strong, structured backgrounds.

Foreground-background cPCA. A significant line of work, initiated by cPCA [1], seeks directions of high variance in a *foreground* dataset that are absent in a separate *background* dataset. This methodology, including subsequent probabilistic and sparse extensions [33, 8, 54, 53], is fundamentally designed for *case–control* settings in which a pure background dataset is explicitly available. Formally, these methods take as input two distinct datasets: a target dataset X, which contains both signal and nuisance variation, and a background-only dataset Y, which contains only nuisance variation. In contrast, our setting is structurally different. The proposed methods (PCA+ and PCA++) address the *positive-pair* scenario, which is central to modern self-supervised learning (e.g., SimCLR [11], MoCo [25]). Instead of relying on a clean background sample, we are given paired observations (X, X^+) that share the same latent signal but are corrupted by independent background variations. The objective is to identify directions that are *shared across positive pairs*, i.e., invariant to nuisance variability, without requiring an explicit background dataset. See additional discussion in Appendix J.

Canonical correlation analysis (CCA). Although there is a conceptual connection between our work and CCA [26], the two approaches differ fundamentally. Both aim to identify shared structure between paired datasets, but they impose different objectives and constraints, which leads to substantial performance differences in noisy, high-dimensional settings. A detailed discussion of the relationship to CCA is provided in Appendix K.

2 Contrastive Factor Model

We observe paired high-dimensional samples $\{(x_i, x_i^+)\}_{i=1}^n$ generated by a **contrastive factor model**, which decomposes each point into a shared signal, an independent background, and noise:

$$x_i = \frac{Aw_i}{Aw_i} + \frac{Bh_i}{Bh_i} + \frac{\varepsilon_i}{E_i}, \qquad x_i^+ = \frac{Aw_i}{Aw_i} + \frac{Bh_i'}{Bh_i'} + \frac{\varepsilon_i'}{E_i'}, \qquad i = 1, \dots, n,$$
 (2.1)

where $A \in \mathbb{R}^{d \times k}$ and $B \in \mathbb{R}^{d \times m}$ have orthonormal columns spanning the signal and background subspaces, respectively; $w_i \in \mathbb{R}^k$ is the *shared* signal factor, while $h_i, h_i' \in \mathbb{R}^m$ are the independent background factors; $\varepsilon_i, \varepsilon_i' \in \mathbb{R}^d$ are independent *random noise*.

Equivalently, we may write

$$A := \left[\sqrt{\lambda_{A,1}} v_{A,1}, \dots, \sqrt{\lambda_{A,k}} v_{A,k} \right], \quad B := \left[\sqrt{\lambda_{B,1}} v_{B,1}, \dots, \sqrt{\lambda_{B,m}} v_{B,m} \right]$$

each column $v_{A,j}$ (resp. $v_{B,j}$) is a principal direction with variance $\lambda_{A,j}$ (resp. $\lambda_{B,j}$). By contrasting each pair (x_i, x_i^+) , which share the same signal Aw_i but differ in background Bh_i , Bh_i' and noise, our goal is to isolate and recover the low-dimensional signal subspace spanned by the columns of A. To make the model analytically tractable, we impose three assumptions:

Assumption 2.1 (Orthogonal signal and background) *The column spaces of the signal and background loading matrices are mutually orthogonal, i.e.,* $span(A) \perp span(B)$.

Assumption 2.2 (Gaussian latent factors) $w_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, I_k), h_i, h_i' \stackrel{\text{iid}}{\sim} \mathcal{N}(0, I_m), all independent across i.$

Assumption 2.3 (Isotropic noise) $\varepsilon_i, \varepsilon_i' \stackrel{\text{iid}}{\sim} \mathcal{N}(0, I_d)$, independent of z_i, w_i, w_i'

On the orthogonality assumption. Orthogonality assumption is made primarily for analytical tractability by forcing the signal and background into disjoint subspaces. Our framework's core mechanism is, however, robust to its violation. The key insight is that even if signal and background subspaces overlap, the *contrastive energy* (as analyzed in Lemma F.1) of the shared directions remains strictly positive, while that of pure background directions is zero. This allows PCA++ to distinguish the full signal space from the background. The main effect of overlap is a reduction in the generalized eigenvalue for the shared directions, but they remain detectable.

On the Gaussian latent factor and noise assumption. This assumption was also made for analytical convenience, as it allows for the clean derivation of exact constants and closed-form error rates. However, we expect the core results to hold more generally for sub-Gaussian distributions. Many of the key results from random matrix theory that we rely on (e.g., in Lemmas G.4 and G.5) have well-known extensions beyond the Gaussian case, often requiring only a few finite moments.

On the isotropic noise assumption. We fix the noise variance to one without loss of generality, since any other scale can be absorbed into the singular values of *A* and *B*.

We discuss the potential to relax these assumptions in Appendix E.

3 PCA+: Contrastive PCA via Alignment Only

A key insight from [50, 12, 39] is that contrastive objectives implicitly promote two complementary geometric forces in representation space: **Alignment** (bringing positive pairs close), and **Uniformity** (evenly spreading all features on the hypersphere). Specifically, for a feature map $f: X \to \mathbb{R}^k$, a general *spectral contrastive learning* objective [24] with uniformity-weight $\tau > 0$ is

$$\mathcal{L}_{sp}(f) = \underset{(x,x^{+})}{\mathbb{E}} \left[\left\| f(x) - f(x^{+}) \right\|_{2}^{2} \right] + \tau \cdot \left\| \underset{x}{\mathbb{E}} \left[f(x) f(x)^{\top} \right] - I_{k} \right\|_{F}.$$
 (3.1)

In this work, we consider a closely related objective, adopting the alignment term from [23]:

$$\mathcal{L}(f) = \underbrace{-\mathbb{E}_{(x,x^{+})} \left[f(x)^{\mathsf{T}} f(x^{+}) \right]}_{alignment} + \tau \cdot \underbrace{\left\| \mathbb{E}_{x} \left[f(x) f(x)^{\mathsf{T}} \right] - I_{k} \right\|_{F}}_{uniformity}. \tag{3.2}$$

Both objectives in (3.1) and (3.2) capture the alignment–uniformity trade-off observed in many popular contrastive learning frameworks [11, 37, 46, 51].

In this section, we "turn off" uniformity ($\tau = 0$) in (3.2) and study a linear encoder that optimizes *only* alignment under our contrastive factor model; the following section then analyzes the effect of enforcing uniformity. We show that an "alignment-only" method (PCA+)-i.e. applying PCA to the contrastive covariance-can successfully recover the signal subspace when background variation is mild. However, in high-dimensional regimes with strong background noise, this pure-alignment approach can fail completely. This failure highlights the necessity of a uniformity-type constraint: by

enforcing feature dispersion, we regain robust signal recovery even when background components overwhelmingly dominate.

To isolate the role of alignment in contrastive learning, we study a linear encoder

$$f(x) = V^{\mathsf{T}} x, \quad V \in \mathbb{R}^{d \times k},$$

and maximize alignment between positive pairs (x_i, x_i^+) without uniformity objective $(\tau = 0)$ in (3.2):

$$\min_{V^{\top}V = I_k} -\frac{1}{n} \sum_{i=1}^n (V^{\top}x_i)^{\top} (V^{\top}x_i^+) \quad \Leftrightarrow \quad \max_{V^{\top}V = I_k} \operatorname{tr} \left(\frac{1}{n} V^{\top} X^{\top} X^+ V\right),$$

where $-\frac{1}{n}\sum_{i=1}^{n}(V^{\top}x_{i})^{\top}(V^{\top}x_{i}^{+})$ is the empirical version of $-\mathbb{E}[(V^{\top}x)^{\top}(V^{\top}x^{+})]$, and $X = (x_{1}, \dots, x_{n})^{\top}$, $X^{+} = (x_{1}^{+}, \dots, x_{n}^{+})^{\top}$ with $X, X^{+} \in \mathbb{R}^{n \times d}$. To prevent representational collapse in such an alignment-focused setup, an orthogonality constraint $V^{\top}V = I_{k}$ is necessary. Note that here $\frac{1}{n}X^{\top}X^{+}$ need not be symmetric, so its eigenvalues can be complex and its eigenvectors need not be orthogonal. A simple fix is to symmetrize

$$S_n^+ = \frac{1}{2n} (X^\top X^+ + X^{+\top} X)$$

which is real-symmetric and shares the same population expectation (shown in Theorem 3.1). Plugging S_n^+ into the alignment objective yields

$$\max_{V^{\top}V=I_{h}} \operatorname{tr}(V^{\top}S_{n}^{+}V), \tag{3.3}$$

which is exactly PCA on the "contrastive" covariance S_n^+ . We call this method PCA+–our baseline that captures pure alignment under an orthogonality constraint.

The estimated signal subspace is taken to be the span of the top k eigenvectors of S_n^+ . We first show that S_n^+ is an unbiased estimator of the true signal covariance:

Theorem 3.1 (Unbiasedness of the contrastive covariance estimator) *Under Assumptions 2.1–2.3, the contrastive covariance estimator* S_n^+ *satisfies* $\mathbb{E}[S_n^+] = AA^\top$.

In other words, in expectation, the contrastive covariance matrix S_n^+ perfectly cancels out both background and noise, recovering the population signal covariance AA^\top .

Having established the behaviour of PCA+ estimator at the population level, we now turn to its finite-sample performance. Our main result is a non-asymptotic upper bound on the subspace estimation error of PCA+, incurred by applying PCA to the contrastive covariance matrix S_n^+ . Denote the true subspace span by columns of A by $\mathcal{U}_A := \text{span}\{v_{A,1},\ldots,v_{A,k}\}$ and let $\widehat{\mathcal{U}}_A$ be our estimator. We measure the estimation error between \mathcal{U}_A and $\widehat{\mathcal{U}}_A$ via their subspace distance based on the *principal angles* between \mathcal{U}_A and $\widehat{\mathcal{U}}_A$. The definition of the subspace distance is provided in Appendix A.

Theorem 3.2 (Finite-sample performance of PCA+) *Suppose Assumptions 2.1–2.3 hold, and that the sample size n obeys:*

$$n \ge \frac{C}{\lambda_{A,h}^2} \left(k \lambda_{A,1}^2 + m \lambda_{B,1}^2 + \max(k, m) \lambda_{A,1} \lambda_{B,1} + d \left(\lambda_{A,1} + \lambda_{B,1} + 1 \right) \right) \log^3(n+d),$$

for a sufficiently large universal C. Then, with probability at least $1 - O((n + d)^{-10})$, the distance between the estimated subspace $\widehat{\mathcal{U}}_A$ and the true subspace \mathcal{U}_A satisfies:

$$\operatorname{dist}(\widehat{\mathcal{U}}_{A}, \mathcal{U}_{A}) \lesssim \frac{1}{\lambda_{A,k}} \left(\lambda_{A,1} \sqrt{\frac{k}{n}} + \lambda_{B,1} \sqrt{\frac{m}{n}} + \sqrt{\lambda_{A,1} \lambda_{B,1}} \sqrt{\frac{\max(k,m)}{n}} + \left(\sqrt{\lambda_{A,1}} + \sqrt{\lambda_{B,1}} + 1 \right) \sqrt{\frac{d}{n}} \right) \log^{1/2}(n+d).$$

This bound shows that, when the smallest signal spike $\lambda_{A,k}$ is large relative to the largest background spikes $\lambda_{B,1}$ and the noise level, PCA+ recovers the signal subspace accurately from finite data. However, if the background strength grows or the eigengap $\lambda_{A,k}$ shrinks, the error can blow up. Indeed, in the extreme one-signal-one-background case below, standard PCA+ fails completely to align with the true signal:

Example 3.3 (One-signal, one-background) Let k = m = 1 in model (2.1) and

$$A := \begin{bmatrix} \sqrt{\lambda_{A,1}} e_1 \end{bmatrix}$$
 and $B := \begin{bmatrix} \sqrt{\lambda_{B,1}} e_2 \end{bmatrix}$,

where $e_1 := (1, 0, \dots, 0)^{\mathsf{T}}$ and $e_2 := (0, 1, 0, \dots, 0)^{\mathsf{T}}$ are the first two standard basis vectors in \mathbb{R}^d .

Theorem 3.4 (Failure under strong background) Under this model with $d/n \to c \in (0, +\infty)$ as $n, d \to +\infty$, the leading PCA+ direction \hat{v}_1 of S_n^+ satisfies

$$\lim_{n,d\to+\infty} (\hat{v}_1^\top e_1)^2 \le 2 \frac{\lambda_{A,1}}{\sqrt{\lambda_{B,1}c}}.$$

This result highlights a limitation of PCA+: when the leading background eigenvalue becomes sufficiently large, $\sqrt{\lambda_{B,1}} \gtrsim \lambda_{A,1}/\sqrt{c}$, the top empirical eigenvector fails to align with the true signal direction, with the alignment remaining bounded away from one. To address this issue, we propose a *constrained* contrastive PCA method in the following section that incorporates uniformity to suppress background interference and achieve robust signal recovery even under strong background conditions.

4 Enforcing Uniformity: PCA++ and Its Analysis

The previous section exposed a flaw of pure-alignment PCA (PCA+): it can no longer recover the true signal when background noise is too strong. We now show that enforcing perfect uniformity—i.e. constraining the projected features to have identity covariance—acts as a powerful regularizer, neutralizing even very strong background.

Concretely, we introduce the *hard-uniformity contrastive PCA* (PCA++):

$$\max_{V \in \mathbb{R}^{d \times k}} \underbrace{\operatorname{tr}(V^{\top} S_n^{+} V)}_{alignment}, \quad \text{s.t.} \quad \underbrace{V^{\top} S_n V = I_k}_{uniformity}, \tag{4.1}$$

where $S_n = \frac{1}{n}X^{\top}X$. The hard constraint $V^{\top}S_nV = I_k$ enforces $\frac{1}{n}\sum_{i=1}^n (Vx_i)(Vx_i)^{\top} = I_k$, i.e. perfect uniformity of the projected features. (3.1).

The PCA++ objective (4.1) can be solved in one shot via a single *generalized eigenvalue decomposition* (see proof in Appendix B.1), making it both conceptually simple and computationally efficient. Specifically, one computes

$$S_n^+ v_j = \lambda_j S_n v_j, \quad \forall \ j \in \{1, \dots, d\}.$$
 (4.2)

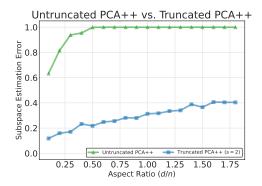
for $\{\lambda_j, v_j\}$, where $S_n^+ = \frac{1}{2n}(X^\top X^+ + X^{+\top} X)$ and $S_n = \frac{1}{n}X^\top X$. The top-k generalized eigenvalues v_1, \ldots, v_k (those with largest λ_j) form the columns of the PCA++ solution $V \in \mathbb{R}^{d \times k}$. When S_n is invertible, (4.2) reduces to the ordinary eigenproblem for $S_n^{-1}S_n^+$. For implementation details see Algorithm 1 in Appendix B.2, and for a full derivation of generalized eigenvalue solvers refer to [38, 19].

In very high dimensions, S_n is often severely ill-conditioned, so solving (4.2) directly may be numerically unstable. Figure 2 (left) shows how–even with the uniformity constraint—the generalized eigenvectors can drift away from the true signal when $d \gg n$. To restore stability, we replace S_n by its rank-s approximation $(S_n)_s := \sum_{j=1}^s \lambda_j v_j v_j^{\mathsf{T}}$ where $\{\lambda_j, v_j\}_{j=1}^s$ are the top-s eigenpairs of S_n , and $s \leq d$ is a tuning parameter. We then solve

$$\max_{V \in \mathbb{R}^{d \times k}} \operatorname{tr}(V^{\top} S_n^+ V), \qquad \text{s.t.} \qquad V^{\top} (S_n)_s V = I_k, \tag{4.3}$$

By construction, $(S_n)_s$ is well-conditioned on its s-dimensional leading subspace and discards the small, unstable directions, enforcing uniformity only where S_n is reliable. As Figure 2 (right) shows, an appropriate s dramatically reduces error, whereas including too many dimensions re-introduces noise. Computationally, one simply replaces S_n with $(S_n)_s$ in Algorithm 1 to obtain a stable solution. We provide detailed guidelines for choosing the value of this crucial hyperparameter in Appendix \mathbb{C} .

Furthermore, this truncation strategy offers a significant computational advantage. Since we only need to find the top-s generalized eigenvectors, we can avoid a full decomposition and instead employ efficient iterative solvers, such as the implicitly restarted Lanczos method (IRLM). This makes the



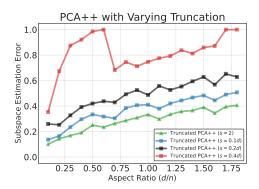


Figure 2: **Effect of covariance truncation on PCA++.** Results are for Example 3.3. **Left:** As d/n increases, truncated PCA++ remains stable and accurate while untruncated PCA++ deteriorates sharply. **Right:** Truncated PCA++ with varying truncation ranks s (fixed s = 2; or s as 0.1d, 0.2d, 0.4d of feature dimension d).

approach scalable to very high-dimensional data, and the details of its computational complexity are discussed in Appendix D.

We will analyze the truncated uniformity-constrained contrastive PCA in two complementary high-dimensional asymptotic regimes using tools from random matrix theory and the spiked covariance model [30, 5, 40]:

- I. **Fixed aspect ratio regime.** The signal eigenvalues $\{\lambda_{A,j}\}_{j=1}^k$ and background eigenvalues $\{\lambda_{B,j}\}_{j=1}^m$ remain constant, while the aspect ratio $d/n \to c > 0$ as both $n, d \to \infty$.
- II. **Growing-spike regime.** Both signal spikes $\lambda_{A,j}$ and background eigenvalues $\lambda_{B,j}$ diverge with $n, d \to \infty$, and the scaled aspect ratios $\frac{d}{n\lambda_{A,j}} \to c_{A,j}$ and $\frac{d}{n\lambda_{B,j}} \to c_{B,j}$, with $0 \le c_{A,1} < \cdots < c_{A,k} < \infty$ and $0 \le c_{B,1} < \cdots < c_{B,m} < \infty$.

Our analysis for the following high-dimensional regimes builds upon the asymptotic alignment between sample and population principal components of $S_n = \frac{1}{n}X^TX$ under model (2.1), as formalized in Lemma G.4. Denote the true population spikes and directions of $\mathbb{E}[xx^T]$ by $\{(\lambda_{A,j}, \nu_{A,j})\}_{j=1}^k$ and $\{(\lambda_{B,j}, \nu_{B,j})\}_{j=1}^m$, sorted in descending order, and let $\{(\hat{\lambda}_{A,j}, \hat{\nu}_{A,j})\}_{j=1}^k$ and $\{(\hat{\lambda}_{B,j}, \hat{\nu}_{B,j})\}_{j=1}^m$ be the corresponding sample eigenpairs obtained by matching each empirical eigenvalue of S_n to its nearest population counterpart in magnitude. This matching preserves the separation between signal and background components in our high-dimensional analysis.

I. Fixed aspect ratio regime. We now study the regime $d/n \to c > 0$ as both $n, d \to \infty$ and ask: when do the sample eigenvalues "stick out" of the Marcenko-Pastur bulk so that both signal and background directions remain identifiable?

Assumption 4.1 (Detectable spikes) The population eigenvalues $\{\lambda_{A,j}\}_{j=1}^k$ and $\{\lambda_{B,j}\}_{j=1}^m$ are all distinct and satisfy $\lambda_{A,j}$, $\lambda_{B,j} \geq \sqrt{c}$ for every j.

Remark. This condition corresponds to the classical BBP threshold, a fundamental requirement in high-dimensional PCA for spiked covariance models. When the signal strength falls below this threshold, the leading eigenvalues of the sample covariance matrix are absorbed into the noise bulk of the Marcenko–Pastur distribution, rendering the signal statistically undetectable [5, 40]. This phenomenon reflects an intrinsic information-theoretic limit of signal detection via PCA in the high-dimensional regime, rather than a limitation specific to our PCA++ method. Under this assumption, the hard-uniformity constrained PCA (PCA++) admits a simple and exact characterization of its asymptotic subspace error:

Theorem 4.2 (Asymptotic subspace error under hard uniformity) *Under Assumptions 2.1–2.3* and 4.1, and $s \ge k$, let $\widehat{\mathcal{U}}_A$ be the top-k subspace returned by the hard-uniformity PCA in (4.3) and

 \mathcal{U}_A be the true population signal subspace. Then as $n, d \to +\infty$ with $n/d \to c \in (0, +\infty)$,

$$\operatorname{dist}(\widehat{\mathcal{U}}_A, \mathcal{U}_A)^2 \longrightarrow 1 - \frac{1 - c\lambda_{A,k}^{-2}}{1 + c\lambda_{A,k}^{-1}} \quad a.s.$$

We can see that when the weakest signal spike $\lambda_{A,k} \gg \sqrt{c}$, the error $1 - \frac{1 - c\lambda_{A,k}^{-2}}{1 + c\lambda_{A,k}^{-1}}$ is small, reflecting accurate recovery; As $\lambda_{A,k} \gg \sqrt{c}$ grows, the error vanishes. Conversely, increasing the aspect ratio c = d/n makes recovery harder and increases the limiting error. Crucially, this result highlights the power of the uniformity constraint. Unlike the pure-alignment method PCA+ (Example 3.3 and Theorem 3.4), which can be overwhelmed by strong background spikes, the hard-uniformity estimator PCA++ continues to recover the signal subspace—albeit with a controlled high-dimensional bias—even when background eigenvalues grow arbitrarily large.

II. Growing-spike regime. In the growing-aspect-ratio regime, where the spikes grow proportionally to d/n such that their "effective signal-to-noise ratios" converge to finite, distinct constants, we no longer require the lower-bound from Assumption 4.1: the growth of the eigenvalues ensures they remain asymptotically separable from the noise bulk.

Assumption 4.3 (Distinct growing spikes) The population eigenvalues $\{\lambda_{A,j}\}_{j=1}^k$ and $\{\lambda_{B,j}\}_{j=1}^m$ are all distinct.

Remark. This assumption was made primarily for analytical convenience. With distinct eigenvalues, we can cleanly map each sample eigenvector to a population counterpart using existing results (Lemmas G.4 and G.5). Without this assumption, if a set of eigenvalues were identical, we would instead estimate a subspace for those directions. Within that estimated subspace, it would be difficult to distinguish which basis vectors correspond to signal and which to background just by looking at S_n alone, making it harder to explain how the uniformity constraint works on a per-direction basis. The key insight is that even if the standard covariance S_n has degenerate subspaces (i.e., multiple identical eigenvalues mixing signal and background components), the contrastive covariance S_n^+ resolves this ambiguity. Since S_n^+ has asymptotically zero energy on pure background directions, the generalized eigenvalue problem can still correctly identify the signal subspace and separate it from the background. We discuss the potential to relax this assumption in Appendix E.3.

The following theorem shows that, the limiting subspace estimation error is a simple function of the weakest signal's effective SNR, $c_{A,k} = \frac{d}{n \lambda_{A,k}}$:

Theorem 4.4 *Under Assumptions* 2.1–2.3 and 4.3, and $s \ge k$, as $n, d \to +\infty$,

$$\operatorname{dist}(\widehat{\mathcal{U}}_A,\mathcal{U}_A)^2 \longrightarrow \frac{c_{A,k}}{1+c_{A,k}} \quad a.s.$$

As $c_{A,k}$ decreases (i.e. the signal grows stronger relative to the ambient dimension d/n), the error vanishes. Moreover, by plugging $c_{A,j} = \frac{d}{n\lambda_{A,j}}$ into the error expression from Theorem 4.2, one recovers exactly the same form in Theorem 4.4, confirming full consistency between the two high-dimensional analyses.

5 Experiments

Simulation studies. We designed a suite of simulations to test our theoretical predictions and to compare standard PCA, the alignment-only method PCA+, our uniformity-constrained PCA++ (PCA++), and the corresponding high-dimensional theory¹. In every scenario, paired samples were drawn from the linear contrastive factor model (Eq. (2.1)), and performance was quantified by the sine of the principal angle between the estimated and true signal subspaces, averaged over fifty independent runs.

We first explored the role of truncation in stabilizing PCA++ when n = 1000 and $\lambda_{A,1} = 10$, $\lambda_{B,1} = 500$. Solving the generalized eigenproblem without truncation leads to erratic signal estimates once d/n

¹The code for reproduction can be found at https://github.com/archer-yang-lab/pcapp.

exceeds about 1 (Figure 2, left). However, projecting the sample covariance onto its top-s subspace with s = 2 fully restores recovery. We then swept s over $\{2, 0.1d, 0.2d, 0.4d\}$ and found that moderate truncation (e.g. s = 0.1d) achieves the best trade-off: it discards the unstable noise directions while preserving enough dimensions to enforce the uniformity constraint (Figure 2, right).

Next, we turned to the asymptotic regimes characterized by Theorems 4.2 and 4.4. Fixing n = 500 and varying the aspect ratio d/n, we embedded a five-dimensional signal subspace (variances [50, 25, 20, 15, 10]) in the first five coordinates and an orthogonal five-dimensional background (variances [500, 400, 300, 200, 100]) in the last five. Applying PCA++ with truncation rank s = 10, we again computed the sine of the largest principal angle to the true signal, averaged over fifty runs.

Under the fixed-aspect-ratio regime, the empirical PCA++ errors trace our closed-form curve almost exactly, while the alignment-only PCA+ method diverges as d/n increases (Figure 3, left). In the growing-spike regime—when both d and spike variances are multiplied by ten—PCA++ continues to adhere to the simple limit $\frac{c}{1+c}$, but PCA+ collapses (Figure 3, right). These results confirm that our theoretical predictions capture the precise high-dimensional behavior of PCA++ in both regimes. Further details and additional simulation results can be found in Appendix H.

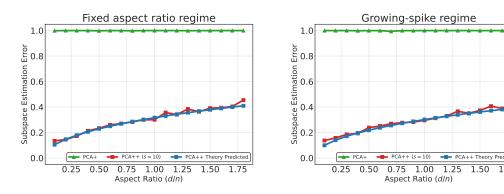


Figure 3: Empirical validation of theoretical predictions for PCA++. Left: Validation in the fixed aspect ratio regime for Theorem 4.2. Right: Validation in the growing-spike regime for Theorem 4.4.

Corrupted MNIST Data. To visually assess signal disentanglement from structured background noise, we created a synthetic dataset of 5,000 paired images by superimposing MNIST digits [16] ('0' or '1') onto distinct ImageNet [15] "grass" patches. Figure 4 compares the 2D embeddings obtained from standard PCA, PCA+, and PCA++. While standard PCA fails to separate classes and PCA+ shows only partial, misaligned separation, PCA++ clearly distinguishes the digits, with separation predominantly along its first learned eigenvector. This visually confirms that the uniformity constraint in PCA++ enables effective isolation of the true digit signal from the background.

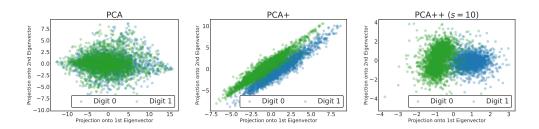


Figure 4: 2D embeddings of noisy digit-over-grass images, standard PCA fails to separate classes. Contrastive PCA+ shows partial, misaligned separation. In contrast, our PCA++ achieves clear class separation predominantly along its first eigenvector, highlighting its superior ability to isolate the true signal and background noise.

Single-cell RNA sequencing data. We evaluated PCA++ on single-cell RNA-seq data from [31], comprising 14,619 control and 14,446 IFN- β -stimulated PBMCs across eight immune cell types. After matching 9,268 cells per condition via donor identity and local structural alignment, we

extracted the top 50 components with both PCA and PCA++ and visualized each embedding (in Figure 5) using UMAP [35]. Standard PCA often separates control and stimulated cells of the same type, even when transcriptional changes are minimal, while PCA++ yields tightly overlapping clusters for invariant populations (e.g., CD4 T cells, B cells, NK cells), while still reflecting biologically meaningful dispersion in responsive cell types (e.g., monocytes in Figure 6). This demonstrates PCA++'s ability to isolate condition-invariant structure in complex single-cell data. Further details and additional simulation results can be found in Appendix I.

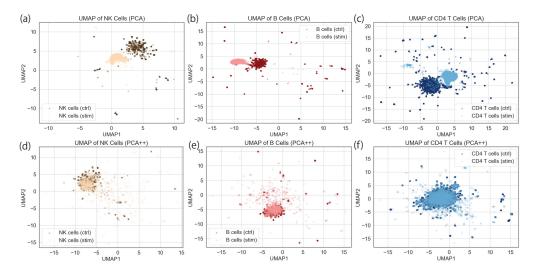


Figure 5: PCA vs. PCA++ embeddings. We apply PCA and PCA++ to matched control and stimulated PBMCs (9,268 cells each) from the [31] dataset and visualize the top 50 components using UMAP. (a–c) show PCA embeddings of CD4 T cells, B cells, and NK cells, where control and stimulated cells are often separated despite minimal biological response. (d–f) show the same cells under PCA++, where alignment across conditions improves, highlighting PCA++'s ability to isolate stable, condition-invariant structure.

6 Discussion and Future Work

This work revisits classical dimensionality reduction through a modern contrastive perspective, showing that alignment alone is insufficient for robust signal recovery in high-dimensional settings with strong structured noise. We demonstrate that incorporating an explicit uniformity constraint—a fundamental ingredient of recent contrastive learning theory—provides a principled safeguard against nuisance variation. The proposed method, PCA++, recovers latent signal subspaces in regimes where both standard PCA and alignment-only contrastive PCA (PCA+) provably fail.

Uniformity as a robustness mechanism. A key conceptual contribution of this work is clarifying the role of uniformity in contrastive learning. Prior studies emphasized its geometric effects (e.g., feature dispersion and isotropy), but its statistical role under structured noise was unclear. Our analysis shows that uniformity suppresses directions aligned with dominant background spikes, acting as a spectral filter that reweights principal directions based on their signal-to-background ratio rather than their raw variance. This connects contrastive learning with classical robust statistics, highlighting that structured interference cannot be eliminated by alignment alone and instead requires principled regularization.

Broader implications and future work. Although studied in a linear setting, the insights extend beyond PCA: (a) *Self-supervised learning*. PCA++ formalizes the benefit of uniformity in SimCLR-style losses and explains why explicit control of feature dispersion improves robustness to nuisance variation. (b) *Multiview learning*. Our paired-data setting parallels CCA, but we show that variance-based alignment alone is insufficient under antagonistic background structure. The success of PCA++ suggests that modern representation learning in multiview settings may benefit from hard covariance constraints. Future extensions include contrastive sparse PCA [55], contrastive kernel PCA [44] and tensor PCA [2] for nonlinear/multiway data, and spectral clustering variants [36] for community detection. These directions promise to expand uniformity-constrained inference across diverse high-dimensional settings.

Impact Statement

This paper aims to advance the field of machine learning. While there may be societal impacts, none require specific attention here.

Acknowledgments

We sincerely thank Yixuan Li for invaluable assistance with the real data analysis. This work was partially supported by NSERC Discovery Grants (RGPIN-2024-06780 and RGPIN-2018-06484) and FRQNT Team Research Project Grant (FRQNT 327788).

References

- [1] Abid, A., Zhang, M. J., Bagaria, V. K., and Zou, J. (2018). Exploring patterns enriched in a dataset with contrastive principal component analysis. *Nature Communications*, 9(1):2134. 1.1
- [2] Anandkumar, A., Ge, R., Hsu, D. J., Kakade, S. M., Telgarsky, M., et al. (2014). Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15(1):2773–2832. 6
- [3] Bachman, P., Hjelm, R. D., and Buchwalter, W. (2019). Learning representations by maximizing mutual information across views. *Advances in Neural Information Processing Systems*, 32. 1.1
- [4] Bai, Z. D. and Yin, Y. Q. (1993). Limit of the smallest eigenvalue of a large dimensional sample covariance matrix. *The Annals of Probability*, 21(3):1275 1294. G
- [5] Baik, J. and Silverstein, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, 97(6):1382–1408. 1, 4, 4, G.4
- [6] Bansal, P., Kavis, A., and Sanghavi, S. (2024). Understanding contrastive learning via Gaussian mixture models. *arXiv preprint arXiv:2411.03517*. 1.1
- [7] Bardes, A., Ponce, J., and LeCun, Y. (2021). VICReg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*. 1.1
- [8] Boileau, P., Hejazi, N. S., and Dudoit, S. (2020). Exploring high-dimensional biological data with sparse contrastive principal component analysis. *Bioinformatics*, 36(11):3422–3430. 1.1
- [9] Campbell, J. Y., Lettau, M., Malkiel, B. G., and Xu, Y. (2001). Have individual stocks become more volatile? an empirical exploration of idiosyncratic risk. *The Journal of Finance*, 56(1):1–43.
- [10] Candès, E. J., Li, X., Ma, Y., and Wright, J. (2011). Robust principal component analysis? *Journal of the ACM*, 58(3):1–37. I.2
- [11] Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pages 1597–1607. PMLR. 1, 1.1, 3
- [12] Chen, T., Luo, C., and Li, L. (2021a). Intriguing properties of contrastive losses. *Advances in Neural Information Processing Systems*, 34:11834–11845. 1.1, 3
- [13] Chen, X. and He, K. (2021). Exploring simple siamese representation learning. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15750–15758.
- [14] Chen, Y., Chi, Y., Fan, J., Ma, C., et al. (2021b). Spectral methods for data science: A statistical perspective. *Foundations and Trends in Machine Learning*, 14(5):566–806. G.1, G.2
- [15] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 248–255. IEEE. 5, I.1

- [16] Deng, L. (2012). The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142. 5, I.1
- [17] Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56. 1
- [18] Fortin, J.-P., Parker, D., Tunç, B., Watanabe, T., Elliott, M. A., Ruparel, K., Roalf, D. R., Satterthwaite, T. D., Gur, R. C., Gur, R. E., et al. (2017). Harmonization of multi-site diffusion tensor imaging data. *Neuroimage*, 161:149–170. 1
- [19] Ghojogh, B., Karray, F., and Crowley, M. (2019). Eigenvalue and generalized eigenvalue problems: Tutorial. *arXiv preprint arXiv:1903.11240.* 4
- [20] Gutmann, M. and Hyvärinen, A. (2010). Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of The International Conference on Artificial Intelligence and Statistics*, pages 297–304. JMLR Workshop and Conference Proceedings.
- [21] Haghverdi, L., Lun, A. T., Morgan, M. D., and Marioni, J. C. (2018). Batch effects in single-cell RNA-sequencing data are corrected by matching mutual nearest neighbors. *Nature Biotechnology*, 36(5):421–427. 1
- [22] HaoChen, J. Z. and Ma, T. (2022). A theoretical study of inductive biases in contrastive learning. arXiv preprint arXiv:2211.14699. 1.1
- [23] HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. (2021). Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34:5000–5011. 1.1, 3
- [24] HaoChen, J. Z., Wei, C., Kumar, A., and Ma, T. (2022). Beyond separability: Analyzing the linear transferability of contrastive representations to related subpopulations. *Advances in Neural Information Processing Systems*, 35:26889–26902. 1.1, 3
- [25] He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738. 1.1
- [26] Hotelling, H. (1936). Relations between two sets of variates. Biometrika, 28(3-4):321–377. 1.1
- [27] Hubert, L. and Arabie, P. (1985). Comparing partitions. *Journal of Classification*, 2(1):193–218. I.2
- [28] Ji, W., Deng, Z., Nakada, R., Zou, J., and Zhang, L. (2023). The power of contrast for feature learning: A theoretical analysis. *Journal of Machine Learning Research*, 24(330):1–78. 1.1
- [29] Johnson, D. D., Hanchi, A. E., and Maddison, C. J. (2023). Contrastive learning can find an optimal basis for approximately view-invariant functions. In *The International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.* OpenReview.net. 1.1
- [30] Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, 29(2):295–327. 4
- [31] Kang, H. M., Subramaniam, M., Targ, S., Nguyen, M., Maliskova, L., McCarthy, E., Wan, E., Wong, S., Byrnes, L., Lanata, C. M., et al. (2018). Multiplexed droplet single-cell RNA-sequencing using natural genetic variation. *Nature Biotechnology*, 36(1):89–94. 5, 5, I.2
- [32] Leek, J. T., Scharpf, R. B., Bravo, H. C., Simcha, D., Langmead, B., Johnson, W. E., Geman, D., Baggerly, K., and Irizarry, R. A. (2010). Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics*, 11(10):733–739. 1
- [33] Li, D., Jones, A., and Engelhardt, B. (2020). Probabilistic contrastive principal component analysis. *arXiv preprint arXiv:2012.07977*. 1.1

- [34] McAllester, D. and Stratos, K. (2020). Formal limitations on the measurement of mutual information. In *International Conference on Artificial Intelligence and Statistics*, pages 875–884. PMLR. 1.1
- [35] McInnes, L., Healy, J., and Melville, J. (2018). UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv* preprint arXiv:1802.03426. 5, 1.2
- [36] Ng, A., Jordan, M., and Weiss, Y. (2001). On spectral clustering: Analysis and an algorithm. *Advances in Neural Information Processing Systems*, 14. 6
- [37] Oord, A. v. d., Li, Y., and Vinyals, O. (2018). Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748.* 1, 1.1, 3
- [38] Parlett, B. N. (1998). The symmetric eigenvalue problem. SIAM. 4
- [39] Parulekar, A., Collins, L., Shanmugam, K., Mokhtari, A., and Shakkottai, S. (2023). Infonce loss provably learns cluster-preserving representations. In Neu, G. and Rosasco, L., editors, *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 1914–1961. PMLR. 1, 1.1, 3
- [40] Paul, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, pages 1617–1642. 4, 4, G.4
- [41] Salloum, R. and Kuo, C.-C. J. (2022). cPCA++: An efficient method for contrastive feature learning. *Pattern Recognition*, 124:108378. J
- [42] Saunshi, N., Ash, J., Goel, S., Misra, D., Zhang, C., Arora, S., Kakade, S., and Krishnamurthy, A. (2022). Understanding contrastive learning requires incorporating inductive biases. In *International Conference on Machine Learning*, pages 19250–19286. PMLR. 1.1
- [43] Saunshi, N., Plevrakis, O., Arora, S., Khodak, M., and Khandeparkar, H. (2019). A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*, pages 5628–5637. PMLR. 1.1
- [44] Schölkopf, B., Smola, A., and Müller, K.-R. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5):1299–1319. 6
- [45] Shen, D., Shen, H., Zhu, H., and Marron, J. (2016). The statistics and mathematics of high dimension low sample size asymptotics. *Statistica Sinica*, 26(4):1747. G.5, G
- [46] Sohn, K. (2016). Improved deep metric learning with multi-class n-pair loss objective. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc. 3
- [47] Tan, Z., Zhang, Y., Yang, J., and Yuan, Y. (2024). Contrastive learning is spectral clustering on similarity graph. In *International Conference on Learning Representations*. 1.1
- [48] Tschannen, M., Djolonga, J., Rubenstein, P. K., Gelly, S., and Lucic, M. (2019). On mutual information maximization for representation learning. *arXiv preprint arXiv:1907.13625*. 1.1
- [49] van der Maaten, L. and Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605. I.2
- [50] Wang, T. and Isola, P. (2020). Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR. 1, 1, 1.1, 3
- [51] Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. (2018). Unsupervised feature learning via non-parametric instance discrimination. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3733–3742. 3
- [52] Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. (2021). Barlow Twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR. 1.1

- [53] Zhang, E. and Li, D. (2025). Contrastive functional principal component analysis. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 39, pages 22380–22388.
- [54] Zhou, Q., Gao, Q., Wang, Q., Yang, M., and Gao, X. (2023). Sparse discriminant pca based on contrastive learning and class-specificity distribution. *Neural Networks*, 167:775–786. 1.1
- [55] Zou, H., Hastie, T., and Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286. 6

Appendix

PCA++: How Uniformity Induces Robustness to Background Noise in Contrastive Learning

A Principal Angles

Denote the true subspace by $\mathcal{U}_A := \operatorname{span}\{v_{A,1}, \dots, v_{A,k}\}$ and let $\widehat{\mathcal{U}}_A$ be an estimator obtained from a given learning method. We measure the estimation error between \mathcal{U}_A and $\widehat{\mathcal{U}}_A$ via their principal angles, which capture the maximal deviations between corresponding directions in the two subspaces:

Definition A.1 (Principal angles) Let \mathcal{U} and \mathcal{U}' be k-dimensional subspaces of \mathbb{R}^d . Let the columns of $U \in \mathbb{R}^{d \times k}$ and $U' \in \mathbb{R}^{d \times k}$ form orthonormal bases for \mathcal{U} and \mathcal{U}' , respectively. Write the singular values of $U^{\top}U'$ in descending order as $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_k \geq 0$. The principal angles θ_j between \mathcal{U} and \mathcal{U}' are defined by

$$\theta_i := \arccos(\sigma_i), \quad for \ j = 1, \dots, k.$$

These angles capture the worst-case misalignment between the two subspaces. A common aggregate distance is

$$\operatorname{dist}_{\sin,\|\cdot\|}(\mathcal{U},\mathcal{U}') := \|\sin\Theta\| := \|\operatorname{diag}(\sin\theta_1,\ldots,\sin\theta_k)\|,$$

where $\|\cdot\|$ is any unitarily invariant matrix norm (e.g., the operator or Frobenius norm).

Remark A.2 This metric is invariant under orthogonal transformations and relates directly to other subspace distances commonly employed in the literature, such as the projection Frobenius norm $||UU^{\top} - U'U'^{\top}||_F$. Thus, it offers an interpretable measure to evaluate subspace estimation accuracy, particularly useful in high-dimensional analysis.

B Generalized Eigenvalue Problem

B.1 Equivalence between (4.1) and (4.2)

Use variable V' instead of V in (4.1) (the reason will become clear later), and introduce a symmetric multiplier $M \in \mathbb{R}^{k \times k}$ for (4.1). The Lagrangian is

$$L(V', M) = \operatorname{tr}(V'^{\mathsf{T}} S_n^{\mathsf{+}} V') - \operatorname{tr}[M(V'^{\mathsf{T}} S_n V' - I_k)].$$

Since M is symmetric, we may rewrite

$$L(V', M) = \operatorname{tr}(V'^{\top} S_n^+ V') - \operatorname{tr}(M V'^{\top} S_n V') + \operatorname{tr}(M).$$

We differentiate L(V', M) w.r.t. V' and then set $\nabla_{V'}L = \mathbf{0}$, we obtain

$$\nabla_{V'}L = 2S_n^+V' - 2S_nV'M = \mathbf{0} \qquad \Longrightarrow \qquad S_n^+V' = S_nV'M. \tag{B.1}$$

Meanwhile, multiply V'^{\top} to both sides of (B.1), we obtain

$$V'^{\top}S_n^+V' = V'^{\top}S_nV'M \Longrightarrow M = V'^{\top}S_n^+V'$$

where we used $V'^{\top}S_nV' = I_k$. The objective function we want to maximize is $\operatorname{tr}(V'^{\top}S_n^+V')$, which is equal to $\operatorname{tr}(M)$. The matrix $M = V'^{\top}S_n^+V'$ is a $k \times k$ real symmetric matrix. Therefore, it is orthogonally diagonalizable. Let Q be a $k \times k$ orthogonal matrix $(Q^{\top}Q = QQ^{\top} = I_k)$ such that $Q^{\top}MQ = \Lambda$, where $\Lambda = \operatorname{diag}(\lambda_1, \lambda_2, \dots, \lambda_k)$ is a diagonal matrix containing the eigenvalues of M.

Define a new matrix V = V'Q. We check if V satisfies the constraint:

$$V^{\mathsf{T}} S_n V = (V'Q)^{\mathsf{T}} S_n (V'Q) = Q^{\mathsf{T}} V'^{\mathsf{T}} S_n V' Q = Q^{\mathsf{T}} I_k Q = Q^{\mathsf{T}} Q = I_k$$

So V also satisfies the constraint. Now substitute $V' = VQ^{T}$ into equation (B.1):

$$S_n^+(VQ^\top) = S_n(VQ^\top)M$$

Post-multiply by Q:

$$S_n^+ V Q^\top Q = S_n V Q^\top M Q.$$

We get

$$S_n^+ V = S_n V \Lambda \tag{B.2}$$

If we write $V = (v_1, v_2, \dots, v_k)$, then equation (B.2) can be expressed column by column as:

$$S_n^+ v_j = \lambda_j S_n v_j$$
 for $j = 1, \dots, k$.

This is precisely the generalized eigenvalue equation. Since the objective function is

$$\operatorname{tr}(V'^{\top}S_n^+V') = \operatorname{tr}(M) = \sum_{i=1}^k \lambda_j$$

To maximize this sum, we must choose the k generalized eigenvectors v_j that correspond to the k largest generalized eigenvalues λ_j . The solution matrix V will have these k eigenvectors as its columns. Typically, the full generalized eigenvalue problem $S_n^+v_j=\lambda_jS_nv_j$ is solved for all d possible eigenvectors and eigenvalues. Then, these eigenvalues are sorted in descending order, $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$, and the eigenvectors v_1, \ldots, v_k corresponding to the k largest eigenvalues are chosen to form the columns of $V=(v_1,\ldots,v_k)$.

B.2 Generalized eigenvalue solver

To compute the hard-uniformity PCA (PCA++) in (4.1), we must solve the generalized eigenvalue problem

$$S_n^+ v_j = \lambda_j S_n v_j, \quad \forall j \in \{1, \dots, d\}.$$

where $S_n^+ = \frac{1}{2n}(X^T X^+ + X^{+T} X)$ and $S_n^- = \frac{1}{n}X^T X$. Equivalently, in matrix form,

$$S_n^+ V = S V \Lambda$$
,

where $V = (v_1, \dots, v_d)$ contains the eigenvectors as columns, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_d)$ is a diagonal matrix of eigenvalues.

A stable and efficient procedure is as follows:

Algorithm 1 Generalized eigenvalue solver

- 1. Compute eigendecomposition: $S_n V_x = V_x \Lambda_x$.
- 2. Compute $R \leftarrow V_x(\Lambda_x + \varepsilon I_d)^{-1/2}$ where ε is a small regularizer to handle nearly zero eigenvalues.
- 3. Project the contrastive covariance: $M \leftarrow R^{\top} S_n^+ R$.
- 4. Compute eigendecomposition: $MU = U\Lambda$.
- 5. Recover generalized eigenvectors $V \leftarrow V_x U$ and generalized eigenvalues Λ .

C On Selecting Hyperparameter s

The choice of *s* in truncated PCA++ controls a fundamental trade-off between the stability of the generalized eigenproblem and the effectiveness of the uniformity constraint:

- A small s ensures that the inverse of the truncated covariance $(S_n)_s$ is well-conditioned, promoting numerical stability. However, it may discard dimensions needed to enforce uniformity effectively, potentially biasing the result.
- A large s enforces uniformity over a larger subspace but risks instability if S_n is ill-conditioned, as including directions with near-zero eigenvalues can amplify noise. As shown in our experiments (Figure 2, left).

Based on this trade-off, we propose to combine the following practical strategies for selecting s:

- Information-based criterion (lower bound): As in standard PCA, one can determine a minimum s by examining the cumulative variance explained by the eigenvalues of the sample covariance S_n . For example, choose s large enough to capture a significant portion (e.g., 90%) of the total variance. This ensures the uniformity constraint operates on the most meaningful directions.
- Stability-based criterion (upper bound): To ensure numerical stability, one can monitor the **condition number** of the truncated matrix $(S_n)_s$, which is the ratio of its largest to its s-th eigenvalue $(\lambda_1(S_n)/\lambda_s(S_n))$. One should choose s such that this ratio remains below a reasonable threshold, avoiding severe ill-conditioning.

In our experiments (e.g., Figure 2, right), we found that PCA++ is robust across a reasonable range of s.

D Computational Complexity

The computational cost of PCA++ has two main components:

- 1. Covariance matrix formation: We compute the sample covariance S_n and the contrastive covariance S_n^+ . For data matrices of size $n \times d$ (samples \times features), forming these $d \times d$ matrices requires matrix multiplications with a complexity of $O(nd^2)$. In the common high-dimensional setting where $d \gg n$, this can be optimized to $O(dn^2)$ by first computing the $n \times n$ Gram matrix.
- 2. **Generalized eigenvalue problem (GEP) solution:** The cost is dominated by the initial truncated eigendecomposition of S_n to rank s. Using an iterative solver like the IRLM, as implemented in scipy, this step has a complexity of approximately $O(sd^2)$. Since the truncation rank s is typically much smaller than d, this step is highly efficient.

Overall, the complexity is comparable to standard PCA, primarily driven by the feature dimension d.

Empirical Scalability To demonstrate its practical performance, we benchmarked the runtime of PCA++ while varying the number of samples (n) and features (d). The results confirm our theoretical analysis.

Table 1: Computationa	l cost of PCA++ (in seconds)
-----------------------	-------------------	-------------

$n \backslash d$	100	1000	5000	10000
100	0.002	0.157	2.288	9.492
1000	0.003	0.204	3.781	15.666
5000	0.009	0.662	13.919	53.610
10000	0.017	1.310	26.433	100.656

Setup: Truncation rank s = 10, top k = 5 eigenvectors estimated. Benchmarked on an Intel Xeon CPU @ 2.20GHz.

The benchmarks show that the runtime is dominated by the feature dimension d, scaling quadratically, while it scales nearly linearly with the number of samples n. This profile makes PCA++ computationally feasible for typical high-dimensional datasets with tens of thousands of features, confirming its practical scalability.

E Discussion of Theoretical Assumptions

E.1 On Orthogonal signal and background (Assumption 2.1)

Why the core mechanism is robust to overlapping subspaces: The key insight, enabled by the linearity of our factor model in (2.1), is that we can analyze the behavior of the covariance matrices in a shared basis. Even with overlap, the contrastive covariance S_n^+ isolates signal components, while the standard covariance S_n accumulates variance from both signal and background in the shared directions.

Let's sketch this for a simple case. Because the model is linear, we can define an orthonormal basis that accounts for the overlap. Suppose span(A) and span(B) share a single direction v_0 . We can decompose the subspaces as:

- **Signal space:** span(A) is spanned by $\{v_0, v_{A,1}, ..., v_{A,k-1}\}$, where v_0 is the shared part and $v_{A,i}$ are the pure signal directions, orthogonal to v_0 .
- **Background space:** span(B) is spanned by { $v_0, v_{B,1}, ..., v_{B,m-1}$ }, where $v_{B,j}$ are the pure background directions.

The population covariances then become:

• Contrastive covariance:

$$E[S_n^+] = \lambda_A v_0 v_0^\top + \sum_{i=1}^{k-1} \lambda_{A,i} v_{A,i} v_{A,i}^\top,$$

Here, $\lambda_{A,0}$ is the signal variance in direction ν_0 . Thus, contributions from the pure background directions $\{\nu_{B,j}\}$ are still averaged out. The resulting expectation remains spanned only by the signal-related directions $\{\nu_0, \nu_{A,1}, ..., \nu_{A,k-1}\}$. The shared direction ν_0 is *not* cancelled.

• Standard covariance:

$$E[S_n] = (\lambda_{A,0} + \lambda_{B,0}) v_1 v_1^{\top} + \sum_{i=1}^{k-1} \lambda_{A,i} v_{A,i} v_{A,i}^{\top} + \sum_{i=1}^{m-1} \lambda_{B,i} v_{B,i} v_{B,i}^{\top}.$$

Here, $\lambda_{B,0}$ is the background variance in direction v_0 . The variance in the shared direction v_0 is amplified by both signal and background components.

Intuition—why PCA++ is robust to violations of orthogonality assumption: The robustness of PCA++ stems from how the generalized eigenvalue problem $S_n^+ v = \lambda S_n v$ interacts with these modified covariance structures. We can analyze this through the lens of the asymptotic contrastive energy for each direction, as defined in our analysis for **Lemma F.1**.

- Pure background directions $(v_{B,j})$: The contrastive energy $v_{B,j}^{\top} S_n^+ v_{B,j}$ remains asymptotically zero, so these directions are filtered out.
- Shared direction (v_0) : The contrastive energy $v_0^\top S_n^+ v_0$ is strictly positive due to the signal component. However, its variance in the standard covariance is now amplified by both signal $(\lambda_{A,0})$ and background $(\lambda_{B,0})$ components.

The resulting contrastive energy for shared direction will be smaller than that of a pure signal spike, but it will still be bounded away from zero. Therefore, the PCA++ objective still detects the shared direction as part of the signal space. The fundamental mechanism—isolating all directions with non-zero contrastive energy—remains intact.

Empirical validation with overlapping subspaces: To provide strong empirical evidence, we ran new simulations with non-orthogonal signal and background subspaces and will add these to the Appendix. We ran new simulations based on the setting in Section H.4, but introduced overlap between the signal and background subspaces. We aligned two background directions with the two weakest signal directions, creating a two-dimensional shared subspace. We tested this under two background noise levels.

- **Signal variances:** [50, 25, 20, 15, 10].
- **Moderate noise background:** [500, 400, 300] (pure background) + [25, 12.5] (shared background).
- **Large noise background:** [500, 400, 300] (pure background) + [100, 50] (shared background).

The results for the fixed aspect ratio regime and growing-spike regime are shown in Table 2–5. For demonstration, we compare against our original theoretical predictions (derived under orthogonality).

Table 2: Fixed			

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++ PCA++ theory		0.207 0.179								

Table 3: Fixed aspect ratio with large overlapping noise

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++	0.254	0.206	0.278	0.280	0.357	0.302	0.359	0.405	0.391	0.416
PCA++ theory	0.104	0.179	0.229	0.268	0.301	0.330	0.356	0.379	0.400	0.410

Table 4: Growing-spike regime with moderate overlapping noise

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++ PCA++ theory	0.121 0.100									

Table 5: Growing-spike regime with large overlapping noise

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++ PCA++ theory					0.316 0.287					

Conclusion: The results in Table 2–5 show that PCA++ is remarkably robust, even when the orthogonality assumption is violated. In the moderate overlapping noise case, the empirical error of PCA++ continues to track the theoretical predictions remarkably well. This demonstrates that when the background variance in the shared subspace is not excessively large, the impact of the overlap is minimal.

In the large overlapping noise case, we observe a slight increase in estimation error, as expected. This is because the background noise in the shared directions becomes strong enough to reduce the effective signal-to-noise ratio, making recovery more challenging. Nevertheless, even in this challenging scenario, PCA++ remains stable and successfully recovers the signal subspace with controlled error, confirming that perfect orthogonality is not a practical prerequisite for our method's success.

E.2 On the Gaussian latent factor assumption (Assumption 2.2):

This assumption was made for analytical convenience, as it allows for the clean derivation of exact constants and closed-form error. However, we expect the core results to hold more generally for sub-Gaussian distributions.

To provide strong empirical evidence for this claim, we have run new simulations where the Gaussian latent factors (w_i, h_i) and noise (ϵ_i) from Assumptions 2.2 and 2.3 were replaced with samples from a standardized Beta(2, 2) distribution (a symmetric, bounded, non-Gaussian distribution). We repeated the experiments from Section H.4 (Figure 3) under this new setting.

The results in Table 6–7 show that the empirical performance of PCA++ under Beta-distributed noise continues to align almost perfectly with our theoretical predictions, which were derived under the Gaussian assumption.

Table 6: Fixed aspect ratio with Beta distribution

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++ PCA++ theory		0.195 0.179								

Table 7: Growing-spike regime with Beta distribution

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++ PCA++ theory			0.223 0.218							

This remarkable consistency provides strong evidence that the Gaussian assumption is a technical choice for analytical clarity rather than a strict requirement for the validity of our results, which appear to exhibit universality.

E.3 On Assumption 4.3 (Distinct growing spikes)

This assumption was made primarily for analytical convenience. Even if the standard covariance S_n has degenerate subspaces (i.e., multiple identical eigenvalues mixing signal and background components), the contrastive covariance S_n^+ resolves this ambiguity. Since S_n^+ has asymptotically zero energy on pure background directions, the generalized eigenvalue problem can still correctly identify the signal subspace and separate it from the background.

To empirically validate this claim, we have run a new set of simulations for the setting in Section H.4 (Figure 3) but with:

- Signal variances: [50, 50, 20, 15, 10].
- Background variances: [500, 500, 300, 50, 50].

The results in Table 8–9 below show that even with these degeneracies, the empirical subspace error for PCA++ continues to track our theoretical predictions, which were derived under the distinct spike assumption. This provides strong evidence that the assumption is a technical convenience rather than a practical necessity.

Table 8: Fixed aspect ratio regime

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++ PCA++ theory						0.323 0.330				

Table 9: Growing-spike regime

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
PCA++	0.136	0.195	0.228	0.270	0.307	0.321	0.353	0.384	0.403	0.384
PCA++ theory	0.100	0.171	0.218	0.256	0.287	0.315	0.339	0.361	0.381	0.391

F Proofs

F.1 Proof of Theorem 3.1

Proof Stacking *n* paired samples into data matrices gives

$$X = W^{\mathsf{T}} A^{\mathsf{T}} + H^{\mathsf{T}} B^{\mathsf{T}} + Z, \qquad X^{\mathsf{T}} = W^{\mathsf{T}} A^{\mathsf{T}} + H'^{\mathsf{T}} B^{\mathsf{T}} + Z',$$
 (F.1)

where:

$$X = (x_1, \dots, x_n)^{\top} \in \mathbb{R}^{n \times d}, \quad X^+ = (x_1^+, \dots, x_n^+)^{\top} \in \mathbb{R}^{n \times d},$$

$$W = (w_1, \dots, w_n) \in \mathbb{R}^{k \times n}, \quad H = (h_1, \dots, h_n) \in \mathbb{R}^{m \times n}, \quad H' = (h'_1, \dots, h'_n) \in \mathbb{R}^{m \times n},$$

$$Z = (\varepsilon_1, \dots, \varepsilon_n)^{\top} \in \mathbb{R}^{n \times d}, \quad Z' = (\varepsilon'_1, \dots, \varepsilon'_n)^{\top} \in \mathbb{R}^{n \times d}.$$

In the contrastive factor model, the covariance of each observation x_i decomposes additively into signal, background, and noise:

$$\mathbb{E}[x_i x_i^\top] = AA^\top + BB^\top + I_d.$$

This decomposition ensures that the spectrum of $\mathbb{E}[x_i x_i^{\mathsf{T}}]$ directly reflects the signal, background, and noise contributions.

Recall that

$$S_n^+ = \frac{1}{2n} (X^\top X^+ + X^{+\top} X)$$

 $S_n^+ = \frac{1}{2n}(X^\top X^+ + X^{+\top} X).$ Substituting the matrix form of the generative model (F.1), we obtain the decomposition

$$S_n^+ = AA^\top + \frac{1}{2} \left(E_n + E_n^\top \right),$$

where the error matrix E_n collects all cross-terms:

$$E_{n} = \frac{1}{n} \left(A(WW^{\top} - nI_{k})A^{\top} + AWH'^{\top}B^{\top} + AWZ' + BHW^{\top}A^{\top} + BHH'^{\top}B^{\top} + BHZ' + Z^{\top}W^{\top}A^{\top} + Z^{\top}H'^{\top}B^{\top} + Z^{\top}Z' \right).$$
(F.2)

By Assumptions 2.1-2.3, each latent factor matrix has independent, zero-mean Gaussian columns, and the noise matrices are independent with zero mean. Consequently,

$$\mathbb{E}[E_n] = \mathbf{0},$$

and hence

$$\mathbb{E}[S_n^+] = AA^\top,$$

as claimed.

F.2 Lemma F.1

Lemma F.1 formalizes the behavior of the contrastive sample covariance matrix S_n^+ when projected onto the sample signal and background eigenvectors derived from S_n :

Lemma F.1 (Contrastive energy of sample directions) Under Assumptions 2.1-2.3 and 4.1, as $n, d \to +\infty$ with $d/n \to c \in (0, +\infty)$, for each $1 \le j \le k$

$$\lim_{n,d\to+\infty}\frac{1}{\hat{\lambda}_{A,j}}\hat{v}_{A,j}^{\top}S_n^+\hat{v}_{A,j}\geq\lim_{n,d\to+\infty}\frac{1}{2}\left(1+\frac{\lambda_{A,j}}{\hat{\lambda}_{A,j}}\frac{1-c\lambda_{A,j}^{-2}}{1+c\lambda_{A,j}^{-1}}-\frac{(1+\sqrt{c})^2}{\hat{\lambda}_{A,j}}\right)\quad a.s.,$$

and for each $1 \le j \le m$

$$\lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{B,i}} \hat{v}_{B,j}^{\top} S_n^{+} \hat{v}_{B,j} = 0 \quad a.s..$$

In other words, in the fixed-c limit the contrastive covariance S_n^+ retains strictly positive "energy" along each true signal direction, but asymptotically vanishes along every background direction.

Proof of Lemma F.1 Proof We begin by noting a useful identity. For any vector $v \in \mathbb{R}^d$

$$v^{\top} S_{n}^{+} v = \frac{1}{2n} v^{\top} \left(X^{\top} X^{+} + X^{+\top} X \right) v,$$

$$= \frac{1}{2n} v^{\top} \left(X^{\top} X^{+} + (X^{\top} X^{+})^{\top} \right) v,$$

$$= \frac{1}{n} v^{\top} X^{\top} X^{+} v,$$
(F.3)

where the last equality follows from the symmetry property of quadratic forms. Specifically, for any matrix $M \in \mathbb{R}^{d \times d}$ and vector $a \in \mathbb{R}^d$:

$$a^{\mathsf{T}} M a = a^{\mathsf{T}} M^{\mathsf{T}} a$$
.

With this identity in hand, our task reduces to proving the following two limits as $n, d \to +\infty$ with

$$\lim_{n,d\to+\infty}\frac{1}{\hat{\lambda}_{A,j}}\hat{v}_{A,j}^{\top}\frac{1}{n}X^{\top}X^{+}\hat{v}_{A,j}\geq\lim_{n,d\to+\infty}\frac{1}{2}\left(1+\frac{\lambda_{A,j}}{\hat{\lambda}_{A,j}}\frac{1-c\lambda_{A,j}^{-2}}{1+c\lambda_{A,j}^{-1}}-\frac{(1+\sqrt{c})^{2}}{\hat{\lambda}_{A,j}}\right),\quad 1\leq j\leq k\quad\text{a.s.},$$

$$\lim_{n,d\to+\infty}\frac{1}{\hat{\lambda}_{B,j}}\hat{v}_{B,j}^{\top}\frac{1}{n}X^{\top}X^{+}\hat{v}_{B,j}=0,\quad 1\leq j\leq m\quad\text{a.s.}$$

Step 1: Proof of $\lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{B,j}} \hat{v}_{B,j}^{\top} \frac{1}{n} X^{\top} X^{+} \hat{v}_{B,j} = 0$

Recall the decomposition

$$X^{\top}X^{+} - nAA^{\top} = nE_{n}$$

$$= A(WW^{\top} - nI_{k})A^{\top} + AWH'^{\top}B^{\top} + AWZ'$$

$$+ BHW^{\top}A^{\top} + BHH'^{\top}B^{\top} + BHZ'$$

$$+ Z^{\top}W^{\top}A^{\top} + Z^{\top}H'^{\top}B^{\top} + Z^{\top}Z'.$$
(F.4)

We analyze each term M in E_n to show

$$\frac{1}{n}\hat{v}_{B,j}^{\top}M\hat{v}_{B,j}\xrightarrow{n\to+\infty}0\quad\text{a.s..}$$

Case 1: $M = nAA^{T}$

Since $A^{\top}\hat{v}_{B,j}$ is the projection of $\hat{v}_{B,j}$ onto the column space of A and $\hat{v}_{B,j}$ itself aligns with the column space o B, Lemma G.4 implies:

$$A^{\top}\hat{v}_{B,j} = \begin{pmatrix} \sqrt{\lambda_{A,1}} v_{A,1}^{\top} \hat{v}_{B,j} \\ \vdots \\ \sqrt{\lambda_{A,k}} v_{A,k}^{\top} \hat{v}_{B,j} \end{pmatrix} \xrightarrow{n \to +\infty} \vec{0} \in \mathbb{R}^{k} \quad \text{a.s..}$$

Hence,

$$\hat{v}_{B,i}^{\top} A A^{\top} \hat{v}_{B,j} \xrightarrow{n \to +\infty} 0$$
 a.s..

Case 2: $M = A(WW^{T} - nI)A^{T}$

Since $WW^{\top} \in \mathbb{R}^{k \times k}$ is of fixed dimension, the strong law of large numbers implies

$$\frac{1}{n}WW^{\top} \xrightarrow{n \to +\infty} I_k \quad \text{a.s..}$$

Hence.

$$\frac{1}{n}\hat{v}_{B,j}^{\top}A(WW^{\top}-nI)A^{\top}\hat{v}_{B,j} \leq \left\|\frac{1}{n}WW^{\top}-I\right\|_{2}\left\|A^{\top}\hat{v}_{B,j}\right\|_{2}^{2} \xrightarrow{n \to +\infty} 0 \quad \text{a.s..}$$

An analogous argument applies to the terms

$$\frac{1}{n}\hat{v}_{B,j}^{\top}AWH^{\top}B^{\top}\hat{v}_{B,j}, \frac{1}{n}\hat{v}_{B,j}^{\top}BH^{+}H^{\top}B^{\top}\hat{v}_{B,j} \text{ and } \frac{1}{n}\hat{v}_{B,j}^{\top}BH^{+}H^{\top}B^{\top}\hat{v}_{B,j}.$$

This shows that all such terms almost surely converge to zero.

Case 3: M = BHZ'

Since Z' and $\hat{v}_{B,j}$ are independent, we can condition on $\hat{v}_{B,j}$ and write

$$Z'\hat{v}_{B,j} = \begin{pmatrix} z_1 \\ \vdots \\ z_n \end{pmatrix},$$

where z_1, \ldots, z_n are i.i.d. Gaussian random variables with mean zero and variance $\|\hat{v}_{B,j}\|_2^2 = 1$. These random variables are independent of $\hat{v}_{B,j}$. Consequently, H is also independent of $Z'\hat{v}_{B,j}$ and given $Z'\hat{v}_{B,j}$, the entries of $HZ'\hat{v}_{B,j}$ are i.i.d. random variables with mean zero and variance $\|Z'\hat{v}_{B,j}\|_2^2$. By the strong law of large numbers,

$$\frac{1}{n^2} ||HZ'\hat{v}_{B,j}||_2^2 = \frac{1}{n^2} (||HZ'\hat{v}_{B,j}||_2^2 - m||Z'\hat{v}_{B,j}||_2^2) + \frac{1}{n^2} (m||Z'\hat{v}_{B,j}||_2^2 - mn) + \frac{mn}{n^2}$$

$$\xrightarrow{n \to +\infty} 0 \quad \text{a.s.}$$

Hence,

$$\frac{1}{n}\hat{v}_{B,j}^{\top}BHZ'\hat{v}_{B,j} \leq \frac{1}{n}||B^{\top}\hat{v}_{B,j}||_2||HZ'\hat{v}_{B,j}||_2 \xrightarrow{n \to +\infty} 0 \quad \text{a.s.},$$

where the limit of $||B^{\top}\hat{v}_{B,j}||_2$ is bounded almost surely by Lemma G.4. A similar argument applies to show

$$\frac{1}{n}\hat{v}_{B,j}^{\top}AWZ'\hat{v}_{B,j} \xrightarrow{n \to +\infty} 0 \quad \text{a.s..}$$

Case 4: $M = Z^T Z'$

Using the same reasoning as in Case 3, the entries of $Z'\hat{v}_{B,j}$ are i.i.d. Gaussian random variable with mean zero and variance 1, and are independent of both $\hat{v}_{B,j}$ and Z. Consequently, conditioning on $Z\hat{v}_{B,j}$, implies that $\frac{1}{n}\hat{v}_{B,j}^{\mathsf{T}}Z^{\mathsf{T}}Z'\hat{v}_{B,j}$ follows a Gaussian distribution with mean zero and variance $\frac{1}{n^2}||Z\hat{v}_{B,j}||_2^2$. By Lemma G.3,

$$\lim_{n,d \to +\infty} \frac{1}{n} ||Z\hat{v}_{B,j}||_2^2 \le (1 + \sqrt{c})^2 < +\infty \quad \text{a.s..}$$

Thus

$$\frac{1}{n}\hat{v}_{B,j}^{\top} Z^{\top} Z' \hat{v}_{B,j} \xrightarrow{n \to +\infty} 0 \quad \text{a.s..}$$

Case 5: $M = Z^T W^T A^T$

From Cases 1, 2 and 4, we already know

$$\begin{split} & ||A^{\top}\hat{v}_{B,j}||_2 \xrightarrow{n \to +\infty} 0 \quad \text{a.s.,} \\ & \frac{1}{n}WW^{\top} \xrightarrow{n \to +\infty} I \quad \text{a.s.,} \\ & \frac{1}{n}||Z\hat{v}_{B,j}||_2^2 \le (1+\sqrt{c})^2 \quad \text{a.s..} \end{split}$$

We now use these facts to bound

$$\begin{split} \frac{1}{n} \hat{v}_{B,j}^\top Z^\top W^\top A^\top \hat{v}_{B,j} &\leq \frac{1}{n} \|Z \hat{v}_{B,j}\|_2 \|W^\top A^\top \hat{v}_{B,j}\|_2 \\ &\leq \frac{1}{n} \|Z \hat{v}_{B,j}\|_2 \|W W^\top \|_2 \|A^\top \hat{v}_{B,j}\|_2 \\ &\xrightarrow{n \to +\infty} 0 \quad \text{a.s..} \end{split}$$

Case 6: $M = Z^T H'^T B^T$

By the same reasoning as in Case 3, the entries of $H'^{\top}B^{\top}\hat{v}_{B,j}$ are i.i.d. Gaussian variable with mean zero and variance $\|B^{\top}\hat{v}_{B,j}\|$. According to Lemma G.4, $\|B^{\top}\hat{v}_{B,j}\|$ remains bounded almost surely. Consequently,

$$\frac{1}{n}\hat{v}_{B,j}^{\top}Z^{\top}H'^{\top}B^{\top}\hat{v}_{B,j} \leq \frac{1}{n}||Z\hat{v}_{B,j}||_2||H'^{\top}B^{\top}\hat{v}_{B,j}||_2 \xrightarrow{n \to +\infty} 0 \quad \text{a.s.},$$

where the convergence follows from the strong law of large numbers and $\lim_{n,d\to+\infty} \frac{1}{n} ||Z\hat{v}_{B,j}||_2^2 \le (1+\sqrt{c})^2$ is shown in Case 4.

Putting these cases together, all contributions of terms in equation (F.4) converge to zero. Dividing by $\hat{\lambda}_{B,j}$ does not affect the limit, since $\hat{\lambda}_{B,j}$ converges to a deterministic value by Lemma G.4. It follows that

$$\frac{1}{\hat{\lambda}_{B,j}} \hat{v}_{B,j}^{\top} \frac{1}{n} X^{\top} X^{+} \hat{v}_{B,j} \xrightarrow{n \to +\infty} 0 \quad \text{a.s.,}$$

completing the proof of Step 1.

Step 2: Proof of $\lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\top} \frac{1}{n} X^{\top} X^{+} \hat{v}_{A,j} \ge \lim_{n,d\to+\infty} \frac{1}{2} \left(1 + \frac{\lambda_{A,j}}{\hat{\lambda}_{A,j}} \frac{1 - c \lambda_{A,j}^{-2}}{1 + c \lambda_{A,j}^{-1}} - \frac{(1 + \sqrt{c})^{2}}{\hat{\lambda}_{A,j}} \right)$

By the same reasoning employed in Step 1, each term M appearing in equation (F.4) satisfies

$$\frac{1}{n}\hat{v}_{A,j}^{\top}M\hat{v}_{A,j} \xrightarrow{n \to +\infty} 0 \quad \text{a.s.},$$

with the exceptions of AA^{T} and $Z^{T}W^{T}A^{T}$.

Case 1: $M = AA^{T}$ We have

$$\hat{v}_{A,j}^{\top}AA^{\top}\hat{v}_{A,j} = \|A^{\top}\hat{v}_{A,j}\|_{2}^{2} = \left\| \begin{pmatrix} \sqrt{\lambda_{A,1}}v_{A,1}^{\top}\hat{v}_{A,j} \\ \vdots \\ \sqrt{\lambda_{A,k}}v_{A,k}^{\top}\hat{v}_{A,j} \end{pmatrix} \right\|_{2}^{2} \xrightarrow{n \to +\infty} \lambda_{A,j} \frac{1 - c\lambda_{A,j}^{-2}}{1 + c\lambda_{A,j}^{-1}} \quad \text{a.s.,}$$

where the limit follows from Lemma G.4.

Case 2: $M = Z^T W^T A^T$ Note that

$$\begin{split} X^\top X &= AWW^\top A^\top + AWH^\top B^\top + AWZ \\ &+ BHW^\top A^\top + BHH^\top B^\top + BHZ \\ &+ Z^\top W^\top A^\top + Z^\top H^\top B^\top + Z^\top Z. \end{split}$$

Following the argument in Step 1 for each of these summands, we obtain

$$\frac{1}{n}\hat{v}_{A,j}^{\top}X^{\top}X\hat{v}_{A,j} - \lambda_{A,j}\frac{1 - c\lambda_{A,j}^{-2}}{1 + c\lambda_{A,j}^{-1}} - 2\frac{1}{n}\hat{v}_{A,j}^{\top}AWZ\hat{v}_{A,j} - \frac{1}{n}\hat{v}_{A,j}^{\top}Z^{\top}Z\hat{v}_{A,j} \xrightarrow{n \to +\infty} 0 \quad \text{a.s..} \quad (F.5)$$

Since $\frac{1}{n}\hat{v}_{A}^{\top} X^{\top}X\hat{v}_{A,j} = \hat{\lambda}_{A,j}$, it follows that

$$\begin{split} \lim_{n,d \to +\infty} \frac{2}{n} \hat{v}_{A,j}^{\top} A W Z \hat{v}_{A,j} &= \lim_{n,d \to +\infty} \hat{\lambda}_{A,j} - \frac{1}{n} \hat{v}_{A,j}^{\top} Z^{\top} Z \hat{v}_{A,j} - \lambda_{A,j} \frac{1 - c \lambda_{A,j}^{-2}}{1 + c \lambda_{A,j}^{-1}} \\ &\geq \lim_{n,d \to +\infty} \hat{\lambda}_{A,j} - (1 + \sqrt{c})^2 - \lambda_{A,j} \frac{1 - c \lambda_{A,j}^{-2}}{1 + c \lambda_{A,j}^{-1}} \quad \text{a.s.,} \end{split}$$

where the last line follows from Lemma G.3.

Combining the contributions from AA^{\top} and $Z^{\top}W^{\top}A^{\top}$ with the negligible effects of all other terms yields the desired lower bound:

$$\lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\top} \frac{1}{n} X^{\top} X^{+} \hat{v}_{A,j} = \lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\top} A A^{\top} \hat{v}_{A,j} + \lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \frac{1}{n} \hat{v}_{A,j}^{\top} A W Z \hat{v}_{A,j}$$
 (F.6)
$$\geq \lim_{n,d\to+\infty} \frac{1}{2} \left(1 + \frac{\lambda_{A,j}}{\hat{\lambda}_{A,j}} \frac{1 - c \lambda_{A,j}^{-2}}{1 + c \lambda_{A,j}^{-1}} - \frac{(1 + \sqrt{c})^{2}}{\hat{\lambda}_{A,j}} \right) \quad \text{a.s..}$$

This completes the proof of Step 2, and hence the theorem.

F.3 Lemma F.2

We extend the fixed-aspect-ratio analysis (Lemma F.1 and Theorem 4.2) to the growing-spike regime.

Lemma F.2 (Contrastive energy in growing-spike regime) *Under Assumptions 2.1-2.3 and 4.3,* as $n, d \rightarrow +\infty$, the contrastive energy along each sample signal direction satisfies

$$\lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\top} S_n^{+} \hat{v}_{A,j} \geq \frac{1}{2(1+c_{A,j})} + \frac{1}{2(1+c_{A,j})^2},$$

for each $1 \le j \le k$, and

$$\lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{B,i}} \hat{v}_{B,j}^{\top} S_n^+ \hat{v}_{B,j} = 0 \quad a.s..$$

for each background index $1 \le j \le m$.

This lemma shows that—even when both dimension and spike strengths grow—the contrastive covariance S_n^+ , still concentrates nonzero contrastive "energy" on every true signal direction while vanishing on all background directions, thus preserving the clean separation needed for accurate subspace recovery under our hard-uniformity constraint.

Proof of Lemma F.2 Proof The proof follows similar lines as that of Theorem F.1, with modifications to account for differences in the current setting. Specifically, we establish the following key results:

1. As
$$n, d \to +\infty$$
, $\frac{1}{\hat{\lambda}_{B,j}} ||A^{\top} \hat{v}_{B,j}||^2 \to 0$ a.s., for all $1 \le j \le m$.

2. As
$$n, d \to +\infty$$
, $\frac{1}{\hat{\lambda}_{A,j}} ||A^{\top} \hat{v}_{A,j}||^2 \to \frac{1}{(1+c_{A,j})^2}$ a.s., for all $1 \le j \le m$.

3. $\lim_{n,d\to+\infty} \frac{1}{n\hat{\lambda}_{B,i}} \lambda_{\max}(Z^{\top}Z) = c_{B,j}$, for $1 \leq j \leq m$. An analogous result holds for $\hat{\lambda}_{A,j}$.

4.
$$\lim_{n,d\to+\infty} \frac{2}{n\hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\mathsf{T}} AWZ \hat{v}_{A,j} \ge \frac{c_{A,j}}{(1+c_{A,j})^2}$$

Using these results along with analogous arguments from the proof of Theorem F.1, we derive the following lower bound:

$$\begin{split} \lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \hat{v}_{A,j}^\top \frac{1}{n} X^\top X^+ \hat{v}_{A,j} &= \lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \hat{v}_{A,j}^\top A A^\top \hat{v}_{A,j} + \lim_{n,d\to+\infty} \frac{1}{\hat{\lambda}_{A,j}} \frac{1}{n} \hat{v}_{A,j}^\top A W Z \hat{v}_{A,j} \\ &\geq \frac{1}{2(1+c_{A,j})} + \frac{1}{2(1+c_{A,j})^2}. \end{split}$$

We now detail the proofs of the above assertions separately.

Case 1: (Analysis of $\frac{1}{\hat{\lambda}_{B,i}} ||A^{\top} \hat{v}_{B,j}||^2$)

We have

$$\frac{1}{\hat{\lambda}_{B,j}} \|A^{\top} \hat{v}_{B,j}\|_{2}^{2} = \left\| \begin{bmatrix} \sqrt{\frac{\lambda_{A,1}}{\hat{\lambda}_{B,j}}} v_{A,1}^{\top} \hat{v}_{B,j} \\ \vdots \\ \sqrt{\frac{\lambda_{A,k}}{\hat{\lambda}_{B,j}}} v_{A,k}^{\top} \hat{v}_{B,j} \end{bmatrix} \right\|_{2}^{2} \xrightarrow{n \to +\infty} 0 \quad \text{a.s.,}$$

where the convergence follows from Lemma G.5 and the observation that

$$\lim_{n,d\to+\infty} \frac{\lambda_{A,i}}{\hat{\lambda}_{B,j}} = \frac{1}{1+c_j} \frac{\lambda_{A,i}}{\lambda_{B,j}} = \frac{1}{1+c_{B,j}} \frac{c_{B,j}}{c_{A,i}} \le \infty.$$

Case 2: (Analysis of $\frac{1}{\hat{\lambda}_{A,i}}||A^{\top}\hat{v}_{A,j}||^2$)

Similarly, using Lemma G.5, we have

$$\frac{1}{\hat{\lambda}_{A,j}} ||A^{\top} \hat{v}_{A,j}||^{2} = \left\| \begin{bmatrix} \sqrt{\frac{\lambda_{A,1}}{\hat{\lambda}_{A,j}}} v_{A,1}^{\top} \hat{v}_{A,j} \\ \vdots \\ \sqrt{\frac{\lambda_{A,k}}{\hat{\lambda}_{A,j}}} v_{A,k}^{\top} \hat{v}_{A,j} \end{bmatrix} \right\|_{2}^{2} \xrightarrow{n \to +\infty} (1 + c_{A,j})^{-2} \quad \text{a.s..}$$

Case 3: (Analysis of $\frac{1}{n\lambda_{B,i}}\lambda_{\max}(Z^{\top}Z)$)

Since Z^TZ and ZZ^T share identical nonzero eigenvalues, Lemma G.3 implies

$$\lim_{n,d\to+\infty} \lambda_{\max}\left(\frac{1}{d}ZZ^{\top}\right) = \lim_{n,d\to+\infty} \left(1 + \sqrt{\frac{n}{d}}\right)^2 = 1.$$

It then follows that

$$\lim_{n,d\to+\infty} \frac{1}{n\hat{\lambda}_{B,j}} \lambda_{\max}(Z^{\top}Z) = \lim_{n,d\to+\infty} \frac{d}{n\hat{\lambda}_{B,j}} \lambda_{\max}\left(\frac{1}{d}Z^{\top}Z\right) = c_{B,j},$$

for $1 \le j \le m$.

Case 4: (Analysis of $\frac{2}{n\hat{\lambda}_{A,j}}\hat{v}_{A,j}^{\mathsf{T}}AWZ\hat{v}_{A,j}$)

By arguments analogous to those in Case 2 of Step 2 in the proof of Lemma F.1, we deduce that

$$\begin{split} \lim_{n,d \to +\infty} \frac{2}{n \hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\top} A W Z \hat{v}_{A,j} &= \lim_{n,d \to +\infty} 1 - \frac{1}{n \hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\top} Z^{\top} Z \hat{v}_{A,j} - \frac{1}{\hat{\lambda}_{A,j}} \hat{v}_{A,j}^{\top} A A^{\top} \hat{v}_{A,j} \\ &\geq 1 - \frac{c_{A,j}}{(1 + c_{A,j})} - \frac{1}{(1 + c_{A,j})^2} \\ &= \frac{c_{A,j}}{(1 + c_{A,j})^2}, \end{split}$$

for $1 \le j \le m$, which establishes the claimed inequality.

The proof is thus complete.

F.4 Proof of Theorem 3.2

Proof We analyze the distance between the estimated signal subspace $\widehat{\mathcal{U}}_A$ and its population counterpart \mathcal{U}_A using the Davis-Kahan $\sin \Theta$ theorem (Lemma G.2). Recall equation (F.2)

$$S_n^+ = AA^\top + \frac{1}{2} \left(E_n + E_n^\top \right),$$

where AA^{\top} represents the signal component and E_n represents the error matrix. Furthermore, the signal component is expressed as:

$$AA^{\top} = \sum_{j=1}^{k} \lambda_{A,j} v_{A,j} v_{A,j}^{\top}.$$

To apply Lemma G.2, we need to bound $||E_n||_2$ such that:

$$||E_n||_2 \le \left(1 - \frac{1}{\sqrt{2}}\right) |\lambda_{A,k}|.$$

Bounding the error matrix E_n . Recall that the error matrix E_n (F.2):

$$E_n = \frac{1}{n} \Big(A(WW^{\top} - nI_k)A^{\top} + AWH'^{\top}B^{\top} + AWZ'$$
$$+ BHW^{\top}A^{\top} + BHH'^{\top}B^{\top} + BHZ'$$
$$+ Z^{\top}W^{\top}A^{\top} + Z^{\top}H'^{\top}B^{\top} + Z^{\top}Z' \Big).$$

We bound each term in E_n using Lemma G.1. For the first term, since the entries in W are i.i.d. standard Gaussian random variables, we have:

$$\left\| A\left(\frac{1}{n}WW^{\top} - I_k\right)A^{\top} \right\| \leq \|A\|_2^2 \left\| \frac{1}{n}WW^{\top} - I_k \right\|_2 \lesssim \lambda_{A,1} \sqrt{\frac{k \log(n+d)}{n}},$$

provided that $n \ge k \log^3(n+d)$, with probability at least $1 - O((n+d)^{-10})$. Applying similar arguments for the remaining terms, we obtain:

$$\begin{split} ||E_n|| &\lesssim \lambda_{A,1} \, \sqrt{\frac{k \log(n+d)}{n}} + \lambda_{B,1} \, \sqrt{\frac{m \log(n+d)}{n}} + \sqrt{\lambda_{A,1} \lambda_{B,1}} \, \sqrt{\frac{\max(k,m) \log(n+d)}{n}} \\ &+ \sqrt{\lambda_{A,1}} \, \sqrt{\frac{d \log(n+d)}{n}} + \sqrt{\lambda_{B,1}} \, \sqrt{\frac{d \log(n+d)}{n}} + \sqrt{\frac{d \log(n+d)}{n}} + \frac{d \log^2(n+d)}{n} \\ &= \left(\lambda_{A,1} \, \sqrt{\frac{k}{n}} + \lambda_{B,1} \, \sqrt{\frac{m}{n}} + \sqrt{\lambda_{A,1} \lambda_{B,1}} \, \sqrt{\frac{\max(k,m)}{n}} \right. \\ &+ \left(\sqrt{\lambda_{A,1}} + \sqrt{\lambda_{B,1}} + 1\right) \, \sqrt{\frac{d}{n}} + \frac{d \log^{3/2}(n+d)}{n} \bigg) \log^{1/2}(n+d). \end{split}$$

Sufficient condition for *n***.** To satisfy the bound $||E_n||_2 \le \left(1 - \frac{1}{\sqrt{2}}\right)|\lambda_{A,k}|$, it is sufficient to assume:

$$n \ge C \left(\frac{1}{\lambda_{A,k}^2} \left(k \lambda_{A,1}^2 + m \lambda_{B,1}^2 + \max(k, m) \lambda_{A,1} \lambda_{B,1} + d \left(\lambda_{A,1} + \lambda_{B,1} + 1 \right) \right) \log(n + d) + \frac{d \log^2(n + d)}{\lambda_{A,k}} + k \log^3(n + d) \right),$$

for some large enough constant C > 0. The condition in the theorem is established by observing that:

$$\frac{d\log^2(n+d)}{\lambda_{A,k}} \le \frac{d\lambda_{A,1}\log^2(n+d)}{\lambda_{A,k}^2}.$$

This allows us to eliminate the second-to-last term in the above expression, and by similar reasoning, the last term can also be omitted.

Final bound on subspace distance. Using Lemma G.2, the distance between the estimated subspace $\widehat{\mathcal{U}}_A$ and the true subspace \mathcal{U}_A satisfies:

$$\operatorname{dist}(\widehat{\mathcal{U}}_A, \mathcal{U}_A) \leq \frac{2\|E_n\|_2}{\lambda (AA^\top)_k - \lambda (AA^\top)_{k+1}} = \frac{2\|E_n\|_2}{\lambda_{A,k}}.$$

Substituting the bound on $||E_n||_2$, we obtain:

$$\begin{split} \operatorname{dist}(\widehat{\mathcal{U}}_{A},\mathcal{U}_{A}) &\lesssim \frac{1}{\lambda_{A,k}} \left(\lambda_{A,1} \, \sqrt{\frac{k}{n}} + \lambda_{B,1} \, \sqrt{\frac{m}{n}} + \sqrt{\lambda_{A,1} \lambda_{B,1}} \, \sqrt{\frac{\max(k,m)}{n}} \right. \\ & + \left(\sqrt{\lambda_{A,1}} + \sqrt{\lambda_{B,1}} + 1 \right) \sqrt{\frac{d}{n}} + \frac{d \log^{3/2}(n+d)}{n} \bigg) \log^{1/2}(n+d) \end{split}$$

where the last term in the parenthesis, $\frac{d \log^{3/2}(n+d)}{n}$, is dominated by the second-to-last term $\sqrt{\frac{\lambda_{A,1}d}{n}}$, under the assumption:

$$n \gtrsim \frac{d\lambda_{A,1}}{\lambda_{A,k}^2} \log^3(n+d) \ge \frac{d}{\lambda_{A,1}} \log^3(n+d).$$

Therefore, the last term in the parenthesis can be omitted, and the final bound becomes:

$$\operatorname{dist}(\widehat{\mathcal{U}}_{A}, \mathcal{U}_{A}) \lesssim \frac{1}{\lambda_{A,k}} \left(\lambda_{A,1} \sqrt{\frac{k}{n}} + \lambda_{B,1} \sqrt{\frac{m}{n}} + \sqrt{\lambda_{A,1} \lambda_{B,1}} \sqrt{\frac{\max(k,m)}{n}} + \left(\sqrt{\lambda_{A,1}} + \sqrt{\lambda_{B,1}} + 1 \right) \sqrt{\frac{d}{n}} \right) \log^{1/2}(n+d).$$

Thus, the theorem is proved.

F.5 Proof of Theorem 3.4

Proof We aim to show that

$$\lim_{n,d\to+\infty}(\hat{v}_1^\top e_1)^2 \leq 2\frac{\lambda_{A,1}}{\sqrt{\lambda_{B,1}c}}.$$

Let $\hat{\lambda}_1$ be the largest eigenvalue of S_n^+ . Then

$$(\hat{v}_1^{\mathsf{T}} e_1)^2 \le \frac{e_1^{\mathsf{T}} S_n^+ e_1}{\hat{\lambda}_1}.$$

Hence, it suffices to establish the following two statements as $n, d \to +\infty$ with $d/n \to c \in (0, +\infty)$

- 1. $\lim_{n,d\to+\infty} e_1^{\mathsf{T}} S_n^+ e_1 = \lambda_{A,1}$,
- 2. $\lim_{n,d\to+\infty} \hat{\lambda}_1 \ge \frac{\sqrt{\lambda_{B,1}c}}{2}$

Once these are established, the desired bound follows immediately.

Step 1: Proof of $\lim_{n,d\to+\infty} e_1^{\mathsf{T}} S_n^{\mathsf{T}} e_1 = \lambda_{A,1}$ Using the identity (F.3), we aim to demonstrate instead:

$$\lim_{n,d\to+\infty}e_1^\top\frac{1}{n}X^\top X^+e_1=\lambda_{A,1}$$

Recall the decomposition (F.4):

$$\begin{split} X^{\top}X^{+} - nAA^{\top} &= nE_{n} \\ &= A(WW^{\top} - nI_{k})A^{\top} + AWH'^{\top}B^{\top} + AWZ' \\ &+ BHW^{\top}A^{\top} + BHH'^{\top}B^{\top} + BHZ' \\ &+ Z^{\top}W^{\top}A^{\top} + Z^{\top}H'^{\top}B^{\top} + Z^{\top}Z'. \end{split}$$

Following a similar approach as in the Step 1 proof of Theorem F.1, we analyze each term in M in E_n . For any matrix M in E_n we demonstrate:

$$\frac{1}{n}e_1^{\mathsf{T}}Me_1 \xrightarrow{n \to +\infty} 0$$
 a.s.,

with the exception of $Z^{T}W^{T}A^{T}$, which is analyzed separately.

Case 1: $M = nAA^{T}$

For this case, we compute:

$$e_1^{\mathsf{T}} A A^{\mathsf{T}} e_1 = \|A^{\mathsf{T}} e_1\|_2^2 = \|\sqrt{\lambda_{A,1}} e_1^{\mathsf{T}} e_1\|_2^2 = \lambda_{A,1}$$

Case 2: $M = Z^T W^T A^T$

Here, we write

$$\frac{1}{n}e_1^\top Z^\top W^\top A^\top e_1 = \frac{\sqrt{\lambda_{A,1}}}{n}WZe_1 = \frac{\sqrt{\lambda_{A,1}}}{n}\sum_{j=1}^n w_j \varepsilon_{j,1} \xrightarrow{n \to +\infty} 0 \quad \text{a.s.,}$$

where $\varepsilon_{j,1}$ represents the first entry of the noise vector ε_j . The convergence to zero follows because w_j and $\varepsilon_{j,1}$ are independent standard Gaussian random variables and the strong law of large numbers.

This completes the proof of Step 1.

Step 2: Proof of $\lim_{n,d\to+\infty} \hat{\lambda}_1 \ge \frac{\sqrt{\lambda_{B,1}c}}{2}$

Let

$$v = \frac{1}{\sqrt{2}}e_2 + \frac{1}{\sqrt{2}} \frac{Z^{\top} H'^{\top}}{\|Z^{\top} H'^{\top}\|_2}.$$

It is sufficient to show that

- 1. $\lim_{n,d\to+\infty} v^{\mathsf{T}} S_n^+ v \ge \frac{\sqrt{\lambda_{B,1}c}}{2}$,
- 2. $\lim_{n \to +\infty} ||v||_2 = 1$.

Using the identity (F.3) and the decomposition (F.4), we proceed by proving that, for each term M in nAA^{\top} and E_n the following holds:

$$\frac{1}{n}v^{\top}Mv \xrightarrow{n \to +\infty} 0 \quad \text{a.s.},$$

with the exception of $Z^{T}H'^{T}B^{T}$, which satisfies

$$\lim_{n,d\to+\infty} \frac{1}{n} v^{\top} Z^{\top} H'^{\top} B^{\top} v \ge \frac{\sqrt{\lambda_{B,1} c}}{2}.$$
 (F.7)

Case 1: $M = AA^{\top}$

From the definition of v, we have

$$\begin{split} A^\top v &= \sqrt{\lambda_{A,1}} e_1^\top v \\ &= \sqrt{\frac{\lambda_{A,1}}{2}} \frac{e_1^\top Z^\top H'^\top}{\|Z^\top H'^\top\|_2} \\ &= \sqrt{\frac{\lambda_{A,1}}{2}} \frac{1}{\|Z^\top H'^\top\|_2} \sum_{j=1}^n h_j' \varepsilon_{j,1}. \end{split}$$

The term $||Z^{T}H'^{T}||_{2}$ satisfies

$$\frac{1}{nd} \|Z^{\top} H'^{\top}\|_{2}^{2} = \frac{1}{d} \sum_{i=1}^{d} \frac{1}{n} \left(\sum_{j=1}^{n} h'_{j} \varepsilon_{j,i} \right)^{2}.$$

Since h'_j and $\varepsilon_{j,i}$ are independent standard Gaussian random variables, the sum $\frac{1}{\sqrt{n}}\sum_{j=1}^n h'_j\varepsilon_{j,i}$ follows a Gaussian distribution with mean zero and variance $\sum_{j=1}^n \frac{1}{n}h'_j^2$. By the strong law of large numbers, this variance converges to 1 almost surely as $n \to +\infty$. Consequently, we have:

$$\frac{1}{nd} \|Z^{\mathsf{T}} H'^{\mathsf{T}}\|_2^2 \xrightarrow{n \to +\infty} 1 \quad \text{a.s..}$$

Hence.

$$v^{\top} A A^{\top} v = \|A^{\top} v\|_{2}^{2} = \frac{\lambda_{A,1}}{2} \frac{1}{\|Z^{\top} H'^{\top}\|_{2}^{2}} \left(\sum_{i=1}^{n} h'_{j} \varepsilon_{j,1} \right)^{2} \xrightarrow{n \to +\infty} 0 \quad \text{a.s..}$$

Case 2: $M = A(WW^{T} - nI)A^{T}$

The arguments for this case are analogous to Case 2 in Step 1 of the proof of Theorem F.1. The same reasoning applies directly to the terms $A(WW^{\top} - nI)A^{\top}$, $\frac{1}{n}\hat{v}_{B,j}^{\top}AWH^{\top}B^{\top}\hat{v}_{B,j}$, $\frac{1}{n}\hat{v}_{B,j}^{\top}BH^{+}H^{\top}B^{\top}\hat{v}_{B,j}$ and $\frac{1}{n}\hat{v}_{B,j}^{\top}BH^{+}H^{\top}B^{\top}\hat{v}_{B,j}$.

Case 3: $M = Z^T Z'$

Since v is independent of Z', the same arguments used in Case 3 and Case 4 of Step 1 in the proof of Theorem F.1 can be applied here. This reasoning extends to the terms $Z^{T}Z'$, BHZ' and AWZ'.

Case 4: $M = Z^T W^T A^T$

From the Case 1, we know that

$$A^{\top}v \xrightarrow{n \to +\infty} 0$$
 a.s..

The rest of the proof follows exactly the reasoning in Case 5 of Step 1 in the proof of Theorem F.1.

Case 5: $M = Z^T H'^T B^T$

We write:

$$\frac{1}{n} v^{\mathsf{T}} Z^{\mathsf{T}} H'^{\mathsf{T}} B^{\mathsf{T}} v = \frac{1}{\sqrt{2}n} e_2^{\mathsf{T}} Z^{\mathsf{T}} H'^{\mathsf{T}} B^{\mathsf{T}} v + \frac{1}{\sqrt{2}n} \frac{(Z^{\mathsf{T}} H'^{\mathsf{T}})^{\mathsf{T}}}{\|Z^{\mathsf{T}} H'^{\mathsf{T}}\|_2} Z^{\mathsf{T}} H'^{\mathsf{T}} B^{\mathsf{T}} v$$

We will show that the first term converges to zero, and the second term satisfies the inequality in (F.7).

Firstly, consider the term $B^{T}v$:

$$B^{\mathsf{T}}v = \sqrt{\lambda_{B,1}}e_{2}^{\mathsf{T}}v$$

$$= \sqrt{\frac{\lambda_{B,1}}{2}}\left(1 + \frac{e_{2}^{\mathsf{T}}Z^{\mathsf{T}}H'^{\mathsf{T}}}{\|Z^{\mathsf{T}}H'^{\mathsf{T}}\|_{2}}\right)$$

$$= \sqrt{\frac{\lambda_{B,1}}{2}}\left(1 + \frac{1}{\|Z^{\mathsf{T}}H'^{\mathsf{T}}\|_{2}}\sum_{j=1}^{n}h'_{j}\varepsilon_{j,2}\right)$$

$$\xrightarrow{n \to +\infty} \sqrt{\frac{\lambda_{B,1}}{2}} \quad \text{a.s.,} \tag{F.8}$$

where the convergence follows from the same reasoning as in Case 1. Hence,

$$\frac{1}{\sqrt{2}n}e_2^{\mathsf{T}}Z^{\mathsf{T}}H'^{\mathsf{T}}B^{\mathsf{T}}v = \frac{1}{\sqrt{2}}B^{\mathsf{T}}v\frac{1}{n}\sum_{i=1}^n h_j'\varepsilon_{j,2} \xrightarrow{n \to +\infty} 0 \quad \text{a.s..}$$

Now, for the second term:

$$\lim_{n,d\to+\infty}\frac{1}{\sqrt{2}n}\frac{(Z^\top H'^\top)^\top}{\|Z^\top H'^\top\|_2}Z^\top H'^\top B^\top v=\lim_{n,d\to+\infty}\frac{\sqrt{\lambda_{B,1}}}{2n}\|Z^\top H'^\top\|_2=\frac{\sqrt{\lambda_{B,1}c}}{2}.$$

It remains to demonstrate that $\lim_{n,d\to+\infty} ||v||_2 = 1$. We have

$$\begin{split} \|v\|_2^2 &= \left(\frac{1}{\sqrt{2}}e_2 + \frac{1}{\sqrt{2}}\frac{Z^{\mathsf{T}}H'^{\mathsf{T}}}{\|Z^{\mathsf{T}}H'^{\mathsf{T}}\|_2}\right)^{\mathsf{T}} \left(\frac{1}{\sqrt{2}}e_2 + \frac{1}{\sqrt{2}}\frac{Z^{\mathsf{T}}H'^{\mathsf{T}}}{\|Z^{\mathsf{T}}H'^{\mathsf{T}}\|_2}\right) \\ &= 1 + \frac{e_2^{\mathsf{T}}Z^{\mathsf{T}}H'^{\mathsf{T}}}{\|Z^{\mathsf{T}}H'^{\mathsf{T}}\|_2} \xrightarrow{n \to +\infty} 1 \quad \text{a.s.,} \end{split}$$

where the convergence follows from the same arguments as in equation (F.8).

This completes the proof of Step 2.

F.6 Proof of Theorem 4.2

Proof We begin by considering the sample covariance matrix S_n and write its eigendecomposition as

$$S_n V_r = V_r \Lambda_r$$

where we order $V_x = (\hat{v}_{A,1}, \dots, \hat{v}_{A,k}, \hat{v}_{B,1}, \dots, \hat{v}_{B,m}, \dots)$ and $\Lambda_x = (\hat{\lambda}_{A,1}, \dots, \hat{\lambda}_{A,k}, \hat{\lambda}_{B,1}, \dots, \hat{\lambda}_{B,m}, 0, \dots)$. Set $\Lambda_x^{-1/2}$ to be the diagonal matrix whose entries are $\hat{\lambda}_k^{-1/2}$ when $\hat{\lambda}_k > 0$ and 0 otherwise, and define the whitened matrix. The transformed matrix M_n can thus be expressed as:

$$M = \Lambda_x^{-1/2} V_x^{\mathsf{T}} S_n^+ V_x \Lambda_x^{-1/2}$$

The matrix is nonzero only within its upper-left $(k + m) \times (k + m)$ submatrix:

 $= \begin{pmatrix} \frac{1}{\lambda_{A,1}} \hat{v}_{A,1}^{\top} S_{n}^{+} \hat{v}_{A,1} & \cdots & \frac{1}{\sqrt{\lambda_{A,1}} \lambda_{A,k}} \hat{v}_{A,1}^{\top} S_{n}^{+} \hat{v}_{A,k} & \frac{1}{\sqrt{\lambda_{A,1}} \lambda_{B,1}} \hat{v}_{A,1}^{\top} S_{n}^{+} \hat{v}_{B,1} & \cdots & \frac{1}{\sqrt{\lambda_{A,1}} \lambda_{B,m}} \hat{v}_{A,1}^{\top} S_{n}^{+} \hat{v}_{B,m} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{1}{\sqrt{\lambda_{A,k}} \lambda_{A,1}} \hat{v}_{A,k}^{\top} S_{n}^{+} \hat{v}_{A,1} & \cdots & \frac{1}{\lambda_{A,k}} \hat{v}_{A,k}^{\top} S_{n}^{+} \hat{v}_{A,k} & \frac{1}{\lambda_{B,1}} \hat{v}_{A,k}^{\top} S_{n}^{+} \hat{v}_{B,1} & \cdots & \frac{1}{\sqrt{\lambda_{A,k}} \lambda_{B,m}} \hat{v}_{A,k}^{\top} S_{n}^{+} \hat{v}_{B,m} \\ \frac{1}{\sqrt{\lambda_{B,1}} \lambda_{A,k}} \hat{v}_{B,1}^{\top} S_{n}^{+} \hat{v}_{A,1} & \cdots & \frac{1}{\sqrt{\lambda_{B,k}} \lambda_{A,k}} \hat{v}_{B,1}^{\top} S_{n}^{+} \hat{v}_{A,k} & \frac{1}{\lambda_{B,1}} \hat{v}_{B,1}^{\top} S_{n}^{+} \hat{v}_{B,1} & \cdots & \frac{1}{\lambda_{B,m}} \hat{v}_{B,n}^{\top} S_{n}^{+} \hat{v}_{B,m} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{1}{\sqrt{\lambda_{B,m}} \lambda_{A,1}} \hat{v}_{B,m}^{\top} S_{n}^{+} \hat{v}_{A,1} & \cdots & \frac{1}{\sqrt{\lambda_{B,m}} \lambda_{A,k}} \hat{v}_{B,m}^{\top} S_{n}^{+} \hat{v}_{A,k} & \frac{1}{\sqrt{\lambda_{B,m}} \lambda_{B,1}} \hat{v}_{B,m}^{\top} S_{n}^{+} \hat{v}_{B,1} & \cdots & \frac{1}{\lambda_{B,m}} \hat{v}_{B,m}^{\top} S_{n}^{+} \hat{v}_{B,m} \end{pmatrix}.$

We now show that, in the limit, this submatrix $M_{(k+m)\times(k+m)}$ vanishes except for the first k diagonal term entries, which remain positive by Lemma F.1.

Entries involving background directions. Using Lemma F.1, it follows directly that cross terms involving eigenvectors from the background subspace vanish asymptotically, for any $1 \le i, j \le m$,

$$\frac{1}{\sqrt{\hat{\lambda}_{B,i}\hat{\lambda}_{B,j}}}\hat{v}_{B,i}^{\top}S_{n}^{+}\hat{v}_{B,j} \leq \frac{1}{\hat{\lambda}_{B,i}}\hat{v}_{B,i}^{\top}S_{n}^{+}\hat{v}_{B,i}\frac{1}{\hat{\lambda}_{B,j}}\hat{v}_{B,j}^{\top}S_{n}^{+}\hat{v}_{B,j} \xrightarrow{a.s.} 0,$$

and similarly, for $1 \le i \le k$, $1 \le j \le m$,

$$\frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{B,j}}}\hat{v}_{A,i}^{\top}S_{n}^{+}\hat{v}_{B,j} \leq \frac{1}{\hat{\lambda}_{A,i}}\hat{v}_{A,i}^{\top}S_{n}^{+}\hat{v}_{A,i}\frac{1}{\hat{\lambda}_{B,j}}\hat{v}_{B,j}^{\top}S_{n}^{+}\hat{v}_{B,j} \xrightarrow{a.s.} 0.$$

Off-diagonal signal-signal entries. Fix $1 \le i < j \le k$. A similar argument as in the derivation of equation (F.6) shows that

$$\lim_{n,d \to +\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \hat{v}_{A,i}^{\top} S_{n}^{+} \hat{v}_{A,j} = \underbrace{\lim_{n,d \to +\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \hat{v}_{A,i}^{\top} AA^{\top} \hat{v}_{A,j}}_{=:T_{1}} + \underbrace{\lim_{n,d \to +\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \frac{1}{2n} \hat{v}_{A,i}^{\top} (AWZ + Z^{\top}W^{\top}A^{\top}) \hat{v}_{A,j}}_{=:T_{2}}.$$

1. T_1 : We have

$$\lim_{n,d\to+\infty}\frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}}\hat{v}_{A,i}^{\top}AA^{\top}\hat{v}_{A,j}=\lim_{n,d\to+\infty}\frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}}\sum_{l=1}^{k}\hat{\lambda}_{A,l}\hat{v}_{A,i}^{\top}v_{A,l}v_{A,l}^{\top}\hat{v}_{A,j}=0,$$

where the last equality follows directly from Lemma G.4.

2. T_2 : Similar reasoning to equation (F.5) yields:

$$\begin{split} \lim_{n,d\to+\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \frac{1}{n} \hat{v}_{A,i}^{\top} (AWZ + Z^{\top}W^{\top}A^{\top}) \hat{v}_{A,j} &= \lim_{n,d\to+\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \hat{v}_{A,i}^{\top} \frac{1}{n} X^{\top}X \hat{v}_{A,j} \\ &- \lim_{n,d\to+\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \hat{v}_{A,i}^{\top} AA^{\top} \hat{v}_{A,j} - \lim_{n,d\to+\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \hat{v}_{A,i}^{\top} \frac{1}{n} Z^{\top}Z \hat{v}_{A,j} \\ &= - \lim_{n,d\to+\infty} \frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}} \hat{v}_{A,i}^{\top} \frac{1}{n} Z^{\top}Z \hat{v}_{A,j}. \end{split}$$

To demonstrate the almost sure convergence of the remaining expression to zero, we apply a leave-one-out decoupling approach:

- For each sample $1 \le l \le n$ delete row l of X and let $\hat{\lambda}_i^{-l}$ be the corresponding eigenvectors. These eigenvectors satisfy independence from the omitted data point x_l and Davis-Kahan perturbation bounds implies $\|\hat{\lambda}_i \hat{\lambda}_i^{-l}\|_2 = O_p(\frac{1}{n})$.
- Express the term as:

$$\hat{v}_{A,i}^{\top} \frac{1}{n} Z^{\top} Z \hat{v}_{A,j} = \frac{1}{n} \sum_{l=1}^{n} (z_{l}^{\top} \hat{v}_{A,i}^{(-l)})^{\top} (z_{l}^{\top} \hat{v}_{A,j}^{(-l)}) + o_{p}(1/\sqrt{n}),$$

where z_l the l-th row of Z, is independent of $\hat{v}_{A,i}^{(-l)}$ and $\hat{v}_{A,j}^{(-l)}$.

• Following the proof of Lemma G.6, $\sum_{l=1}^{n} (z_l^{\mathsf{T}} \hat{v}_{A,i}^{(-l)})^{\mathsf{T}} (z_l^{\mathsf{T}} \hat{v}_{A,j}^{(-l)})/n$ is an average of i.i.d. zero-mean variables with finite variance, hence it converges to zero almost surely by the strong law of large numbers.

Consequently, we establish:

$$\lim_{n,d\to+\infty}\frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}}\frac{1}{n}\hat{v}_{A,i}^{\top}(AWZ+Z^{\top}W^{\top}A^{\top})\hat{v}_{A,j}=-\lim_{n,d\to+\infty}\frac{1}{\sqrt{\hat{\lambda}_{A,i}\hat{\lambda}_{A,j}}}\hat{v}_{A,i}^{\top}\frac{1}{n}Z^{\top}Z\hat{v}_{A,j}=0,$$

proving that the off-diagonal entries vanish asymptotically.

Positive definite submatrix $M_{n,k \times k}$ We have demonstrated that the asymptotic limit of the matrix $M_{n,(k+m)\times(k+m)}$ except for its first k diagonal term, which remain strictly positive by Lemma F.1. Consequently, asymptotically, exactly k sample eigenvectors correspond to positive eigenvalues, yielding the estimated signal set:

$$V = (\hat{v}_{A,1}, \dots, \hat{v}_{A,k}).$$

Evaluate the estimated subspace Finally, we quantify the accuracy of the estimated subspace $\widehat{\mathcal{U}}_A$ relative to the true signal subspace \mathcal{U}_A . By Lemma G.4, it follows that:

$$\lim_{n,d\to+\infty} V^{\mathsf{T}} A = \lim_{n,d\to+\infty} \begin{pmatrix} \hat{v}_{A,1}^{\mathsf{T}} v_{A,1} & \cdots & \hat{v}_{A,1}^{\mathsf{T}} v_{A,k} \\ \vdots & & \vdots \\ \hat{v}_{A,k}^{\mathsf{T}} v_{A,1} & \cdots & \hat{v}_{A,k}^{\mathsf{T}} v_{A,k} \end{pmatrix}$$
$$= \operatorname{diag} \left(\sqrt{\frac{1 - c\lambda_1^{-2}}{1 + c\lambda_1^{-1}}}, \dots, \sqrt{\frac{1 - c\lambda_k^{-2}}{1 + c\lambda_k^{-1}}} \right)$$

Hence, the asymptotic squared subspace distance is given by:

$$\begin{split} \lim_{n,d\to+\infty} \mathrm{dist}(\widehat{\mathcal{U}}_A,\mathcal{U}_A)^2 &= \lim_{n,d\to+\infty} \max_{1\leq j\leq k} \sin^2(\arccos\sigma_j) \\ &= \lim_{n,d\to+\infty} \max_{1\leq j\leq k} 1 - \sigma_j^2 \\ &= 1 - \frac{1-c\lambda_k^{-2}}{1+c\lambda_k^{-1}}. \end{split}$$

F.7 Proof of Theorem 4.4

Proof The proof follows directly from the arguments detailed in the proof of Theorem 4.2, in conjunction with Lemma G.5 and Lemma F.2. Accordingly, the algorithm asymptotically yields:

$$V = (\hat{v}_{A,1}, \dots, \hat{v}_{A,k}).$$

Consequently, we have:

$$\lim_{n,d\to+\infty} V^{\top} A = \operatorname{diag}\left((1 + c_{A,1})^{-1/2} (1 + c_{A,k})^{-1/2} \right)$$

and the squared subspace distance asymptotically satisfies:

$$\begin{split} \lim_{n,d\to+\infty} \operatorname{dist}(\widehat{\mathcal{U}}_A, \mathcal{U}_A)^2 &= \lim_{n,d\to+\infty} \max_{1 \leq j \leq k} \sin^2(\arccos \sigma_j) \\ &= \lim_{n,d\to+\infty} \max_{1 \leq j \leq k} 1 - \sigma_j^2 \\ &= 1 - \frac{1}{1 + c_{A,k}} \\ &= \frac{c_{A,k}}{1 + c_{A,k}}. \end{split}$$

G Preliminary Lemmas

This section collects technical tools that will be invoked repeatedly in the proofs of our main results.

Lemma G.1 (Lemma 3.5 in [14]) Assume that the entries in matrices $W \in \mathbb{R}^{k \times n}$ and $Z \in \mathbb{R}^{d \times n}$ are i.i.d. standard Gaussian random variables, if $n \geq k \log^3(n+d)$, with probability at least $1 - O((n+d)^{-10})$ one has

$$\begin{split} & \frac{1}{n} \left\| W Z^{\top} \right\| \lesssim \sqrt{\frac{d \log(n+d)}{n}}, \\ & \left\| \frac{1}{n} W W^{\top} - I_k \right\| \lesssim \sqrt{\frac{k \log(n+d)}{n}}, \\ & \left\| \frac{1}{n} Z Z^{\top} - I_d \right\| \lesssim \sqrt{\frac{d \log(n+d)}{n}} + \frac{d \log^2(n+d)}{n}. \end{split}$$

The next lemma is a convenient corollary of the classical Davis-Kahan $\sin \Theta$ theorem.

Lemma G.2 (Corollary 2.8 in [14]) Let M^* and $M = M^* + E$ be two $n \times n$ real symmetric matrices. We express the eigendecomposition of M^* and M as follows

$$M^* = \sum_{i=1}^n \lambda_j^* v_j^* v_j^{*\top} = \begin{pmatrix} U^* & U_{\perp}^* \end{pmatrix} \begin{pmatrix} \Lambda^* & 0 \\ 0 & \Lambda_{\perp}^* \end{pmatrix} \begin{pmatrix} U^{*\top} \\ U_{\perp}^{*\top} \end{pmatrix}.$$

Here, $\{\lambda_j\}_{j=1}^d$ (resp. $\{\lambda_j^*\}_{j=1}^d$) denote the eigenvalues of M (resp. M^*), and v_j (resp. v_j^*) stands for the eigenvector associated with the eigenvalue λ_j (resp. λ_j^*). Additionally, we take

$$\Lambda := \operatorname{diag}([\lambda_1, \cdots, \lambda_r]) \in \mathbb{R}^{r \times r}, \qquad U := [v_1, \cdots, v_r] \in \mathbb{R}^{n \times r}.$$

The matrices U^* , U_{\perp}^* , Λ^* , and Λ_{\perp}^* are defined analogously. Suppose that $|\lambda_1^*| \geq |\lambda_2^*| \geq \cdots \geq |\lambda_r^*| > |\lambda_{r+1}^*| \geq \cdots |\lambda_n^*|$ and $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_n|$ (i.e., the eigenvalues are sorted by their magnitudes). If $||E|| < (1 - 1/\sqrt{2})(|\lambda_r^*| - |\lambda_{r+1}^*|)$, then

$$\mathrm{dist}(\mathcal{U}, \mathcal{U}^*) \leq \sqrt{2} \|\sin \Theta\| \leq \frac{2 \|EU^*\|}{|\lambda_r^*| - |\lambda_{r+1}^*|} \leq \frac{2 \|E\|}{|\lambda_r^*| - |\lambda_{r+1}^*|},$$

where \mathcal{U} (resp. \mathcal{U}^*) is the subspace spanned by U (resp. U^*).

By combining Theorem 2 and Remark 1 of [4], we obtain the following lemma.

Lemma G.3 Assume that the feature vector x has i.i.d. entries with zero mean, unit variance, and bounded 4-th moment. As $n, d \to \infty$, $0 \le c < \infty$,

$$\lim_{n,d \to +\infty} \lambda_{\min}^+(X^\top X/n) = (1 - \sqrt{c})^2 \quad \text{a.s.}$$

$$\lim_{n,d \to +\infty} \lambda_{\max}(X^\top X/n) = (1 + \sqrt{c})^2 \quad \text{a.s.}$$

where λ_{\min}^+ denotes the smallest positive eigenvalue.

Lemma G.4 ([5] and [40]) Under model (2.1), let $\lambda_1 > \lambda_2 > ... > \lambda_{k+m} > \sqrt{c} \in (0, +\infty)$ denote the top k+m eigenvalues. Then, as $d/n \to c$ with $n, d \to +\infty$,

$$\begin{split} \hat{\lambda}_j &\to 1 + \lambda_j + c \frac{1 + \lambda_j}{\lambda_j} \quad a.s., \\ a^\top \hat{v}_j \hat{v}_j^\top b &\to \frac{1 - c \lambda_j^{-2}}{1 + c \lambda_j^{-1}} a^\top v_j v_j^\top b \quad a.s., \end{split}$$

for any deterministic unit vectors $a, b \in \mathbb{R}^d$ and $1 \le j \le k + m$.

Lemma G.5 (Theorem 5.1 in [45]) *Under model* (2.1), *let* $\lambda_1 > \lambda_2 > ... > \lambda_{k+m}$ *denote the top* k+m *eigenvalues. Then, as* $n,d \to +\infty$ *with* $\frac{d}{n\lambda_j} \to c_j \in (0,+\infty)$ *and* $\lambda_{k+m} \to \infty$,

$$\begin{split} \hat{\lambda}_j/\lambda_j &\to (1+c_j) \quad a.s., \\ v_i^\top \hat{v}_j &\to (1+c_j)^{-1/2} \mathbb{1}(j=i) \quad a.s., \end{split}$$

for $1 \le i, j \le k + m$.

Proof The original statement of Theorem 5.1 in [45] covers explicitly the results of both the sample eigenvalues and the alignment between the j-th sample and population eigenvectors when i = j. Here, we extend the arguments to handle the off-diagonal scenario $i \neq j$. Let

$$\hat{\mathbf{v}}_i = (\hat{\mathbf{v}}_{1,i}, \dots, \hat{\mathbf{v}}_{d,i})^{\mathsf{T}}, \quad j = 1, \dots, d$$

denote the eigenvectors of the sample covariance. Without loss of generality, we adopt the standard Euclidean basis $\{e_j\}_{j=1}^d$ for the population eigenvectors, setting $v_j = e_j$. Under this choice, it suffices to demonstrate that, as $n, d \to +\infty$,

$$\hat{v}_{i,j} \to 0$$
 a.s.,

for all $i \neq j$. We first consider the scenario with j = 1. By Theorem 5.1 (equation (8.21)) of [45], it follows directly that:

$$\sum_{i=2}^{k+m} \hat{v}_{i,1}^2 \xrightarrow{n \to +\infty} 0 \quad \text{a.s..}$$

Hence, we conclude $\hat{v}_{i,1} \to 0$ a.s. for all $i \neq 1$, proving the claim for j = 1.

Next, for the case, j = 2, equations (8.23) and (8.24) in Theorem 5.1 from [45] imply, as $n, d \rightarrow +\infty$

$$\sum_{i=2}^{k+m} \hat{v}_{i,2}^2 \to \frac{1}{1+c_2} \quad \text{a.s.,} \quad \sum_{i=2}^{k+m} \hat{v}_{2,i}^2 \to \frac{1}{1+c_2} \quad \text{a.s..}$$

Consequently, we have $\hat{v}_{i,2} \to 0$ a.s. for all $i \neq 2$. Furthermore, from equation (8.17) in [45], we have:

$$\sum_{i=i}^{k+m} \hat{v}_{1,i}^2 \xrightarrow{n \to +\infty} \frac{1}{1+c_1} \quad \text{a.s..}$$

which directly implies $\hat{v}_{1,2} \to 0$ a.s. as $n, d \to +\infty$. Thus, the result also holds for j = 2.

The above arguments can be iteratively applied to subsequent indices j = 3, 4, ..., k + m by analogous reasoning. Hence, we obtain the general result for arbitrary i and j, completing the proof.

Lemma G.6 (Asymptotic orthogonality of Gaussian projections) *Let* $a, b \in \mathbb{R}^d$ *be unit vectors satisfying* $a^{\mathsf{T}}b = 0$, *and let* $X \in \mathbb{R}^{n \times d}$ *be a random matrix with entries* $X_{ij} \overset{i.i.d.}{\sim} \mathcal{N}(0, 1)$. *Then:*

$$\frac{1}{n}a^{\mathsf{T}}X^{\mathsf{T}}Xb \xrightarrow{a.s.} 0 \quad as \ n, d \to +\infty.$$

Proof We first rewrite the expression as:

$$\frac{1}{n}a^{\top}X^{\top}Xb = \frac{1}{n}(Xa)^{\top}(Xb) = \frac{1}{n}\sum_{i=1}^{n}(Xa)_{i}(Xb)_{j}.$$

Let $z_i := (Xa)_i(Xb)_i$, so that:

$$\frac{1}{n}a^{\mathsf{T}}X^{\mathsf{T}}Xb = \frac{1}{n}\sum_{i=1}^{n}z_{j}.$$

Each row $X_{j\bullet} \in \mathbb{R}^d$ of X is an i.i.d. sample from $\mathcal{N}(0, I_d)$. Then:

$$(Xa)_j = X_{j \bullet}^{\mathsf{T}} a, \quad (Xb)_j = X_{j \bullet}^{\mathsf{T}} b,$$

and thus:

$$z_i = (X_{i \bullet}^\top a)(X_{i \bullet}^\top b).$$

Since $X_{j\bullet} \sim \mathcal{N}(0, I_d)$, and $a, b \in \mathbb{R}^d$ are fixed vectors, the random variables $X_{j\bullet}^{\mathsf{T}}a$ and $X_{j\bullet}^{\mathsf{T}}b$ are jointly Gaussian with:

$$\mathbb{E}[X_{j\bullet}^{\mathsf{T}}a] = 0, \quad \mathbb{E}[X_{j\bullet}^{\mathsf{T}}b] = 0, \quad \text{and} \quad \mathsf{Cov}(X_{j\bullet}^{\mathsf{T}}a, X_{j\bullet}^{\mathsf{T}}b) = a^{\mathsf{T}}b = 0.$$

Hence, $X_{i\bullet}^{\top}a$ and $X_{i\bullet}^{\top}b$ are uncorrelated zero-mean Gaussian random variables and thus independent.

Therefore, since both are independent and centered,

$$\mathbb{E}[z_i] = \mathbb{E}[(X_{i \bullet}^\top a)(X_{i \bullet}^\top b)] = \mathbb{E}[X_{i \bullet}^\top a] \cdot \mathbb{E}[X_{i \bullet}^\top b] = 0.$$

Also, z_1, \ldots, z_n are i.i.d. with finite variance. Then by the strong law of large numbers,

$$\frac{1}{n}\sum_{i=1}^n z_i \xrightarrow{\text{a.s.}} \mathbb{E}[z_i] = 0.$$

Thus,

$$\frac{1}{n}a^{\top}X^{\top}Xb \xrightarrow{\text{a.s.}} 0.$$

H Additional Details for Numerical Simulations (Section 5)

Below we describe in additional details of the synthetic data configurations and evaluation procedures underlying our numerical experiments. In all experiments we measure the subspace estimation error as the sine of the principal angle between the estimated and true signal spaces, averaged over 50 Monte Carlo trials.

H.1 One-signal-one background-model-varying background strength

We generated n=2000 paired samples in d=800 dimensions from the linear contrastive factor model (Eq. (2.1)). The true signal subspace is one-dimensional (k=1): $A=[\sqrt{\lambda_{A,1}},0,\ldots,0]^{\top}$ places a spike of magnitude $\lambda_{A,1}=10$ in coordinate 1. In $B=[0,\sqrt{\lambda_{B,1}},0,\ldots,0]^{\top}$, background variation is confined to the orthogonal coordinate 2 (m=1), with strength $\lambda_{B,1}$ varied so that $\lambda_{A,1}/\sqrt{\lambda_{B,1}}$ ranges over $\{0.3125,\ldots,0.666\}$. For each method (PCA, PCA+, PCA++), we compute the sine of the principal angle between the estimated vector and the true signal direction (the first standard basis vector e_1). A fixed truncation rank s=2 is used to stabilize PCA++ (though truncation does not affect this one-dimensional signal case). Results in Figure 1 (left) show that, as background strength grows, PCA collapses to noise, PCA+ eventually drifts onto the background axis, while PCA++ remains tightly aligned with the true signal.

H.2 One-signal-one background-model-varying aspect ratios

We fixed n = 500 and varied d so that d/n spans $\{0.1, \ldots, 1.8\}$. Both signal and background spikes remain at $\lambda_{A,1} = 10$ in coordinate 1 and $\lambda_{B,1} = 500$ in coordinate 2. We apply ordinary PCA, PCA+, and PCA++ with truncation s = 2. The sine of the principal angle to e_1 is plotted versus d/n. Figure 1 (right) shows that while PCA and PCA+ errors grow unbounded for large d/n, PCA++ error remains uniformly low.

H.3 Stability of PCA++: truncation effects

Figure 2 (left) illustrates the instability of untruncated PCA++ for n = 1000 as d/n increases from 0.1 to 1.8. With the same one-dimensional signal ($\lambda_{A,1} = 10$) and background ($\lambda_{B,1} = 500$), direct solution of the generalized eigenproblem yields erratic directions at high aspect ratios. Replacing the sample covariance S_n by its rank-s approximation with s = 2 restores stable recovery.

Figure 2 (right) examines PCA++'s sensitivity to s. Again with n = 1000, $\lambda_{A,1} = 10$, $\lambda_{B,1} = 500$, and $d/n \in [0.1, 1.8]$, we vary $s \in \{2, 0.1d, 0.2d, 0.4d\}$. Moderate truncation (e.g. s = 0.1d) best balances bias and variance, discarding noisy directions while retaining the subspace needed to enforce uniformity.

H.4 Empirical high-dimensional asymptotics

We fix n = 500 and let d/n vary for the fixed-aspect-ratio regime (Theorem 4.2), using a fivedimensional signal k = 5. The signal loading matrix $A \in \mathbb{R}^{d \times 5}$ was defined such that the signal occupied the first five feature dimensions:

$$A = \begin{bmatrix} \sqrt{\lambda_{A,1}} & 0 & 0 & 0 & 0 \\ 0 & \sqrt{\lambda_{A,2}} & 0 & 0 & 0 \\ 0 & 0 & \sqrt{\lambda_{A,3}} & 0 & 0 \\ 0 & 0 & 0 & \sqrt{\lambda_{A,4}} & 0 \\ 0 & 0 & 0 & 0 & \sqrt{\lambda_{A,5}} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

with the signal component variance per dimension $[\lambda_{A,1}, \lambda_{A,2}, \lambda_{A,3}, \lambda_{A,4}, \lambda_{A,5}] = [50, 25, 20, 15, 10]$. The true signal subspace \mathcal{U}_A is spanned by $[e_1, \dots, e_5]$. The background loading matrix $B \in \mathbb{R}^{d \times 5}$ was defined orthogonally to A, affecting the last five feature dimensions:

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \sqrt{\lambda_{B,4}} & 0 \\ 0 & 0 & \sqrt{\lambda_{B,3}} & 0 & 0 \\ 0 & \sqrt{\lambda_{B,2}} & 0 & 0 & 0 \\ \sqrt{\lambda_{B,1}} & 0 & 0 & 0 & 0 \end{bmatrix},$$

with the signal component variance per dimension $[\lambda_{B,1}, \lambda_{B,2}, \lambda_{B,3}, \lambda_{B,4}, \lambda_{B,5}] = [500, 400, 300, 200, 100]$. Truncation rank is s = 10. For the growing-spike regime (Theorem 4.4), we scale both d and all spikes by a factor of 10. As shown in Figure 3, the empirical subspace error of PCA++ tracks the theoretical prediction almost exactly, even as the spikes and dimension diverge, whereas PCA+ consistently fails to align with the true signal.

H.5 Empirical validation with fixed aspect ratio

To assess performance under a fixed aspect ratio of d/n = 0.4, we conducted simulations comparing standard PCA, PCA+, PCA++, and the theoretical prediction. Data were generated from the linear contrastive factor model (Eq. (2.1)) with signal and background each occupying five orthogonal directions in a *p*-dimensional space. The signal subspace consisted of the first five canonical axes, with variances $[\lambda_{A,1}, \lambda_{A,2}, \lambda_{A,3}, \lambda_{A,4}, \lambda_{A,5}] = [20, 20, 15, 10, 10]$, while the background spanned the last five

axes with variances $[\lambda_{B,1}, \lambda_{B,2}, \lambda_{B,3}, \lambda_{B,4}, \lambda_{B,5}] = [500, 500, 200, 100, 100]$ in the "large noise" regime and [100, 100, 50, 25, 25] in the "mild noise" regime. We varied the sample size $n \in \{100, 500, 5000\}$, keeping d = 0.4 n, and performed 50 independent trials for each setting. Subspace error was measured by the largest principal angle between the estimated and true signal subspaces.

Under large background noise (Table 10), both PCA and PCA+ fail completely, yielding error near one, whereas PCA++ achieves low error (0.304 at n = 100, dropping to 0.212 at n = 5000) with small variability (standard deviation from 0.037 to 0.005). Even when background noise is mild (Table 11), PCA+ only improves its average error at large n (0.222 at n = 5000) but remains highly unstable (sd = 0.062). In contrast, PCA++ matches or betters PCA+'s mean error (0.212) while reducing its standard deviation to 0.004. These results demonstrate that PCA++ not only lowers estimation error but also yields far more consistent recovery of the signal subspace than either PCA or the alignment-only method.

Table 10: Fixed Aspect Ratio (d/n = 0.4) under Large Background Noise

Setting	PCA	PCA+	PCA++	PCA++ theory
n = 100, p = 40	0.999 (0)	0.999 (0)	0.304 (0.037)	0.205
n = 500, p = 200	0.999 (0)	0.999 (0)	0.225 (0.010)	0.205
n = 5000, p = 2000	0.999 (0)	0.999 (0.001)	0.212 (0.005)	0.205

Table 11: Fixed Aspect Ratio (d/n = 0.4) under Mild Background Noise

Setting	PCA	PCA+	PCA++	PCA++ theory
n = 100, p = 40	0.999(0)	0.996 (0.011)	0.293 (0.038)	0.205
n = 500, p = 200	0.999(0)	0.713 (0.287)	0.227 (0.010)	0.205
n = 5000, p = 2000	0.999(0)	0.222 (0.062)	0.212 (0.004)	0.205

I Additional Details for Real Data Applications (Section 5)

I.1 Corrupted MNIST data

This section details the experimental setup for the qualitative evaluation on noisy MNIST digits, the results of which are presented in Figure 4. The aim was to visually assess signal disentanglement from structured background noise using a dataset of n = 5000 paired synthetic images (x_i, x_i^+) .

Data generation. The signal components were images of digits '0' and '1' from the MNIST dataset [16]. Each original 28×28 grayscale MNIST image had its pixel values normalized to the range [0, 1]; let M_i denote such a processed digit. Backgrounds were sourced from the "grass" synset of the ImageNet dataset [15]. Full ImageNet images were converted to grayscale, from which 28×28 patches were randomly cropped. Each cropped grass patch was then also normalized so its pixel values ranged from [0, 1]. Let G_i be independently sampled and processed grass patches. Paired images were then constructed through pixel-wise addition:

$$x_i = M_i$$

$$x_i^+ = 0.25 \cdot M_i + G_i.$$

This setup uses the same MNIST digit M_i for both x_i and its positive pair x_i^+ , but with different grass backgrounds, and with the MNIST signal scaled down by a factor of 0.25 in the x_i^+ view. Each resulting synthetic image was flattened into a p = 784 dimensional vector. Finally, before applying the PCA methods, the set of all x_i vectors (forming matrix X) was globally mean-centered (i.e. the mean vector of X was subtracted from each row x_i). Similarly, the set of all x_i^+ vectors (forming matrix X^+) was globally mean-centered.

I.2 Single-cell RNA sequencing data

We investigate the impact of immune stimulation on peripheral blood mononuclear cells (PBMCs) using single-cell RNA sequencing data from [31]. This dataset includes scRNA-seq profiles from

8 human donors collected under two conditions: control and interferon- β (IFN- β) stimulation. The processed data comprise 14,619 control cells and 14,446 stimulated cells, spanning 8 curated immune cell types. These include both IFN- β -responsive populations (e.g., monocytes) and cell types with minimal transcriptional response (e.g., B cells), providing a natural testbed for studying condition-invariant structure. While cells are not explicitly paired in the raw data, shared donor identities and consistent annotations allow post hoc matching for comparative analysis. Each cell is labeled with a curated cell type, offering a stable reference to assess shifts in population structure under perturbation. Our goal is to demonstrate that the proposed PCA++ method more effectively captures biologically meaningful, condition-invariant structure in real-world single-cell data. To this end, we apply both PCA and PCA++, and visualize the resulting embeddings by applying UMAP [35] to the top 50 components of each projection. We compare the structures side-by-side across control and stimulated conditions to assess the preservation and alignment of cell-type clusters.

To prepare the dataset for contrastive analysis, we first filtered the gene space to retain only informative and biologically relevant features. We excluded non-coding genes (e.g., lncRNA, miRNA, pseudogenes) by annotating Ensembl gene IDs using the mygene API and removing entries with non-protein-coding biotypes. We then selected the top 2,000 highly variable genes (HVGs) using the Seurat v3 method implemented in Scanpy. This yielded a final feature space shared across 29,065 cells, comprising 14,619 control and 14,446 stimulated cells.

Because cells are not explicitly paired in the original dataset, we constructed cell pairs post hoc using donor identity and local structural alignment. For each donor, we performed independent normalization, log-transformation, and PCA on cells from the control and stimulated conditions. We clustered cells within each condition using Leiden clustering and computed the centroids of each cluster in the 50-dimensional PCA space. Clusters were matched across conditions by cosine similarity of their centroids. Within each matched pair of clusters, we applied nearest-neighbor matching in PCA space to identify at most 300 well-aligned cell pairs. This strategy ensured both inter-condition comparability and intra-donor consistency, resulting in 9,268 matched control and 9,268 matched stimulated cells across all 8 donors.

To evaluate the stability of cellular identity across conditions, we apply PCA and our proposed PCA++ to the full set of matched control and stimulated cells. Both methods produce 50-dimensional latent embeddings, which we visualize using UMAP. In the PCA++ framework, the stimulated cells are treated as the target dataset and the control cells as the target-plus dataset, enabling the method to extract components that emphasize shared structure while suppressing condition-specific variation. This design aligns with our goal of capturing invariant cell populations under perturbation.

The left panel of Figure 6 shows the UMAP projection of PCA embeddings. While cell types are well-separated, the embeddings exhibit clear batch effects—control and stimulated cells of the same type often appear as disconnected clusters. This suggests that PCA captures both cell identity and stimulation-driven variance. In contrast, the right panel of Figure 6 shows UMAP on PCA++ embeddings. Here, stimulated and control cells are more co-localized, especially for stable populations like CD4 T cells and B cells, indicating that PCA++ better preserves condition-invariant structure. Although some separation remains for highly responsive populations (e.g., monocytes), PCA++ enables clearer alignment of shared cellular manifolds across conditions.

To better interpret the results, we highlight key observations from Figure 5, which contrasts UMAP visualizations derived from PCA (a, b, c) and PCA++ (d, e, f) projections. In the PCA embeddings, cell types are well separated in both control and stimulated conditions; however, control and stimulated cells of the same type frequently occupy non-overlapping regions. This indicates that standard PCA captures both cell identity and condition-induced variation, leading to separation even within the same cell type. In contrast, PCA++ (d, e, f) produces embeddings where control and stimulated cells are better aligned within each cell type. Although the overall separation between cell types becomes slightly less distinct, the coherence of shared structure across conditions improves markedly. For example, invariant populations such as CD4 T cells, B cells, and NK cells remain tightly clustered regardless of stimulation, highlighting PCA++'s ability to emphasize stable transcriptional programs. Importantly, these findings are consistent with known biology: these cell types are relatively insensitive to IFN- β stimulation and exhibit minimal transcriptional changes. In contrast, more responsive cell types such as monocytes appear more dispersed, reflecting biologically meaningful variation rather than projection-induced noise. These results demonstrate that PCA++ more effectively

disentangles cell identity from stimulation effects, enabling clearer interpretation of shared cellular structure across perturbed and unperturbed states.

we have also included a comprehensive comparison across several methods, including standard PCA, alignment-only PCA+ (which lacks the uniformity constraint), PCA++, as well as additional baseline dimensionality reduction techniques (UMAP [35], t-SNE [49], Robust PCA [10]). All methods were evaluated on the same single-cell RNA-seq data, focusing on their ability to recover cell-type groupings in both control and stimulated conditions.

The adjusted Rand index [27, ARI] (lower is better in this context) for each method is summarized in Table 12.

	Cells	B Cells	NK Cells	All
PCA	0.2084	0.2977	0.2502	0.1478
PCA+	0.0036	0.0242	0.0299	0.014
PCA++	0.0008	0.005	0.0049	0.0172
U-MAP	0.2031	0.2506	0.2078	0.1392
t-SNE	0.0233	0.0967	0.0477	0.0217
Robust PCA	0.0399	0.1907	0.1314	0.0443

As shown, both PCA+ and PCA++ dramatically reduce the ARI compared to standard PCA and other popular methods, indicating much better mixing of the two conditions within cell types. Importantly, adding the uniformity constraint in PCA++ further improves the results over PCA+, especially in B cells and NK cells. These quantitative results highlight the utility of both the alignment and uniformity components. PCA++ achieves superior separation of invariant populations compared to alternative approaches.

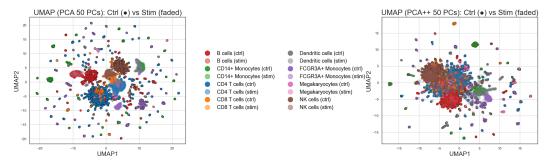


Figure 6: UMAP visualizations of PCA (left) vs. PCA++ (right) on all cells. Both plots show UMAP projections of the top 50 components applied to all matched PBMCs. The PCA embedding (left) separates control and stimulated cells—even within the same cell type—due to confounding condition-specific variance. In contrast, the PCA++ embedding (right), computed using stimulated cells as the target and control cells as auxiliary input, better aligns the two conditions. Invariant populations such as CD4 T cells and B cells remain coherent across stimulation, while more responsive types (e.g., monocytes) exhibit increased spread, reflecting true biological variability.

J Additional Discussion on cPCA

Although there is a fundamental distinction between cPCA and PCA++, one reviewer suggested that cPCA could nonetheless be adapted to address the problem solved by PCA++. Specifically, the reviewer proposed constructing a synthetic foreground and background via $x_f = (x+x^+)/2$ and $x_b = (x-x^+)/2$, and then applying cPCA to the pair (x_f, x_b) . While this construction yields a cPCA-like method that is theoretically valid at the population level, our analysis and experiments demonstrate that it is substantially less stable and reliable in finite-sample regimes compared to PCA++.

We can see that at the population level, the signal is perfectly isolated. Under the contrastive factor model, the population covariances of the synthesized data are:

$$\Sigma_f = \mathbb{E}[x_f x_f^\top] = AA^\top + \frac{1}{2}BB^\top + \frac{1}{2}I_d$$
$$\Sigma_b = \mathbb{E}[x_b x_b^\top] = \frac{1}{2}BB^\top + \frac{1}{2}I_d$$

Therefore, the difference $\Sigma_f - \Sigma_b = AA^{\mathsf{T}}$ perfectly recovers the signal covariance. The practical challenge, however, arises from finite-sample estimation error. Define foreground matrix $X_f = (X + X^+)/2$ and background matrix $X_b = (X - X^+)/2$. The sample covariances of the synthesized data are $\widehat{\Sigma}_f = X_f^{\mathsf{T}} X_f/n$ and $\widehat{\Sigma}_b = X_b^{\mathsf{T}} X_b/n$, respectively.

Regarding cPCA: The reviewer's proposed approach relies on the matrix subtraction $\hat{\Sigma}_f - \alpha \hat{\Sigma}_b$, this subtraction of two large, noisy matrices can amplify noise, potentially overwhelming the true signal and leading to unstable eigenvectors.

Regarding cPCA++: A cPCA-like extension, referred to as cPCA++ [41], has also been proposed in prior work. This method identifies contrastive directions by solving a generalized eigenvalue problem of the form $\widehat{\Sigma}_b^{-1}\widehat{\Sigma}_f$, where $\widehat{\Sigma}_f$ and $\widehat{\Sigma}_b$ are the empirical covariance matrices of the foreground and background data, respectively. As we demonstrate in our paper (Figure 2), directly inverting $\widehat{\Sigma}_b$ is numerically unstable when dimension d is large. The reviewer might suggest stabilizing this by using a truncated pseudoinverse, similar to our PCA++. However, this reveals a deeper conceptual issue: this procedure is equivalent to projecting the foreground data X_f onto the principal subspace of the background data X_b and then performing PCA. Since the synthesized background X_b contains no signal A by construction, its principal subspace is also signal-free. Projecting the foreground onto this signal-free subspace would annihilate the very signal we aim to recover. Therefore, a truncated cPCA++ would fundamentally fail.

Empirical validation:

To investigate these different approaches empirically, we performed a new set of experiments comparing PCA++ against alternative methods proposed by the reviewer.

- 1. **Baselines on synthesized data:** Following the reviewer's suggestion, we evaluated methods that operate on synthesized foreground $X_f = (X + X^+)/2$ and background $X_b = (X X^+)/2$ data. This includes standard cPCA and cPCA++ (based on a difference of covariances). This allows us to directly test the stability of the "subtract-then-decompose" approach versus our "decompose-from-cross-covariance" method.
- 2. **CCA:** We also introduced CCA as a canonical and highly relevant baseline. We apply it by treating the paired data matrices, X and X^+ , as the two views. In our model, the shared signal is the sole source of population-level correlation between these views. Therefore, CCA, which finds directions of maximal correlation, is theoretically suited for recovering the signal subspace and serves as a strong benchmark.

We consider the following experiment setup:

• Experimental setup: The experiment follows the n = 500, k = 5 fixed aspect ratio setting from Sec E.4. For 'cPCA', we set $\alpha = 1$. This is the most principled choice as it yields an unbiased population estimator

$$\mathbb{E}[\hat{\Sigma}_f - \hat{\Sigma}_b] = AA^{\top} + BB^{\top}/2 + I_d/2 - BB^{\top}/2 - I_d/2 = AA^{\top}.$$

We tested two scenarios: one with moderate background noise and one with strong background noise.

- Signal variances: [50, 25, 20, 15, 10].
- Moderate noise background: [100, 50, 40, 30, 20].
- Large noise background: [500, 400, 300, 200, 100].

Experimental results (subspace error):

The results are provided in the following Table 13 and 14. As the results clearly show, the cPCA adaptation is highly unstable even with its optimal hyperparameter setting $\alpha=1$. This, combined with the conceptual issue in the truncated cPCA++ approach, demonstrates that these methods are not well-suited for this problem. In contrast, PCA++ method remains stable, highlighting the robustness of our GEP formulation.

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
cPCA	0.293	0.819	0.430	0.977	0.284	0.626	0.813	0.581	1.000	0.964
cPCA++	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.999	0.999	1.000
CCA	0.160	0.488	0.999	1.000	1.000	1.000	1.000	0.999	1.000	1.000
PCA++	0.125	0.212	0.250	0.275	0.311	0.347	0.375	0.388	0.415	0.401
PCA++ theory	0.104	0.179	0.229	0.268	0.301	0.330	0.356	0.379	0.400	0.410

Table 13: Performance comparison across different aspect ratios with moderate noise background.

Aspect Ratio	0.1	0.3	0.5	0.7	0.9	1.1	1.3	1.5	1.7	1.8
cPCA	1.000	1.000	0.999	0.941	0.998	0.999	0.998	1.000	0.998	0.996
cPCA++	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
CCA	0.154	0.369	0.999	1.000	1.000	1.000	1.000	1.000	1.000	1.000
PCA++	0.156	0.185	0.230	0.283	0.321	0.339	0.379	0.353	0.410	0.447
PCA++ theory	0.104	0.179	0.229	0.268	0.301	0.330	0.356	0.379	0.400	0.410

Table 14: Performance comparison across different aspect ratios with large noise background.

K Additional discussion on the relationship to CCA

While both CCA and PCA++ methods aim to find shared structure between paired datasets, they differ crucially in their objectives and constraints, leading to significant performance gaps in noisy, high-dimensional settings.

Objective functions and constraints:

• Standard CCA: Seeks two distinct projection matrices, U and V, that maximize the correlation between the projected views $U^{T}X$ and $V^{T}X^{+}$. The objective for the leading component is:

$$\max_{u,v} \frac{u^\top (X^\top X^+) v}{\sqrt{(u^\top (X^\top X) u)(v^\top (X^{+T} X^+) v)}}$$

The key instability lies in the denominator, which normalizes by variance from both views $(X \text{ and } X^+)$. In our problem setting, both views contain large, independent background and noise components, making this normalization highly susceptible to noise amplification. Although one can also revise CCA for the positive-pair setting by imposing constrain U = V, the objective becomes a Generalized Rayleigh Quotient, which, to our knowledge, does not have the same closed-form solution as standard CCA.

- PCA++: Our formulation makes two critical changes that enhance stability
 - Shared projection space: It enforces U = V, searching for a single, shared subspace that captures the signal common to both views.
 - **Robust normalization:** Instead of normalizing by the noisy variance from both views, PCA++ maximizes the shared covariance $v^{\mathsf{T}}(X^{\mathsf{T}}X^{\mathsf{T}})v$ subject to a hard uniformity constraint on only one view $v^{\mathsf{T}}(X^{\mathsf{T}}X)v = 1$. This constraint acts as a powerful regularizer, anchoring the solution to a more stable variance structure and preventing distortion from the background and noise present in the second view X^{T} .

Empirical Performance:

This difference in formulation has a dramatic impact on performance. As our new experiments (provided in Table 13 and 14 of Appendix J) demonstrate:

- **Standard CCA is highly unstable** in the settings we study. It fails to recover the signal, with its performance collapsing as dimensionality or noise increases.
- PCA++ remains stable and robust, successfully recovering the signal subspace across all tested regimes.

This confirms that our specific choice of a shared projection space and a single-view uniformity constraint is essential for reliable signal recovery in this problem setting.

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: We support each point by either theory or experiment.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: The paper relies on certain assumptions—such as a linear model structure and noiseless or Gaussian data. It discusses the limitations arising from these assumptions and how results might be affected when they are violated.

Guidelines:

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

3. Theory assumptions and proofs

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

Answer: [Yes]

Justification: See Appendix F and G.

Guidelines:

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

4. Experimental result reproducibility

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

Answer: [Yes]

Justification: See Appendix H.

Guidelines:

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

5. Open access to data and code

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

Answer: [Yes]

Justification: The code and data are publicly released at the time of submission. The experiments are based on standard datasets and use well-established implementations as described in the supplementary material.

Guidelines:

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

6. Experimental setting/details

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

Answer: [Yes]

Justification: See Appendix H

Guidelines:

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

7. Experiment statistical significance

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

Answer: [Yes]

Justification: The differences are often significantly large.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

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

8. Experiments compute resources

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

Answer: [No]

Justification: The experiments do not involve heavy computation, and thus detailed computational resources are not necessary for reproducibility. Relevant implementation details can be found in Appendices B.2 and H.

Guidelines:

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

9. Code of ethics

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

Answer: [Yes]

Justification: We obey all aspects of the Code of Ethics.

Guidelines:

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

10. Broader impacts

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

Answer: [Yes]

Justification: This is a theory-oriented paper with no societal impact.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

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

11. Safeguards

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

Answer: [NA]

Justification: This paper does not involve such models.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [NA]

Justification: The paper does not use existing assets.

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

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

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

Answer: [NA]

Justification: We introduce no new assets.

Guidelines:

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

14. Crowdsourcing and research with human subjects

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

Answer: [NA]

Justification: We do not use crowdsourcing.

Guidelines:

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

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

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

Answer: [NA]

Justification: We do not have such studies.

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

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.