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

Mingqi Wu[1,2], Qiang Sun[3,4], Archer Y. Yang[1,2]

1 McGill University; 2 Mila; 3 University of Toronto; 4 MBZUAI

NEURAL INFORMATION PROCESSING SYSTEMS

1. PCA++ Disentangles Signal from Noise

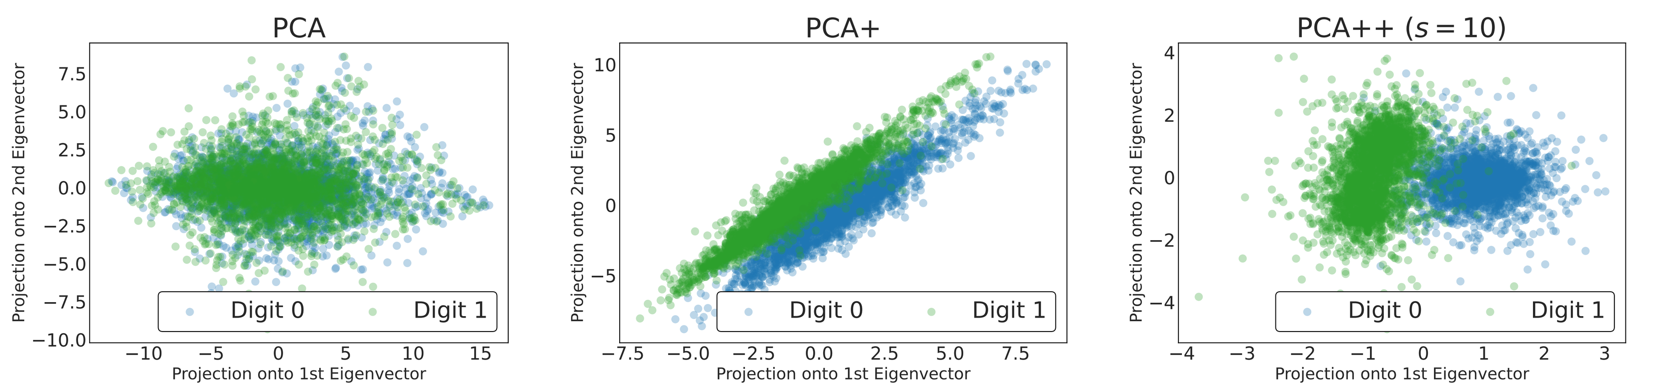

Evidence: PCA++ is the only method that clearly isolates the true digit signal.

2. Problem Setup: The Contrastive Factor Model

We model our data as arriving in positive pairs (x, x^+). Each pair shares a common low-dimensional signal (A) but is corrupted by independent, structured background noise (B).

Our linear factor model is defined as:

$$x_i = Aw_i + Bh_i + \varepsilon_i, \qquad A := \left[\sqrt{\lambda_{A,1}}\lambda_{A,1}, \dots, \sqrt{\lambda_{A,k}}v_{A,k}\right],$$
$$x_i^+ = Aw_i + Bh_i' + \varepsilon_i', \qquad B := \left[\sqrt{\lambda_{B,1}}\lambda_{B,1}, \dots, \sqrt{\lambda_{B,m}}v_{B,m}\right].$$

Our Goal: To recover the shared signal subspace spanned by A.

3. The Problem: Alignment-Only Methods Fail

Standard PCA fails to separate signals from structured background noise. While contrastive learning offers a solution by aligning positive pairs (X, X^+), a simple alignment-only approach (PCA+) is not robust.

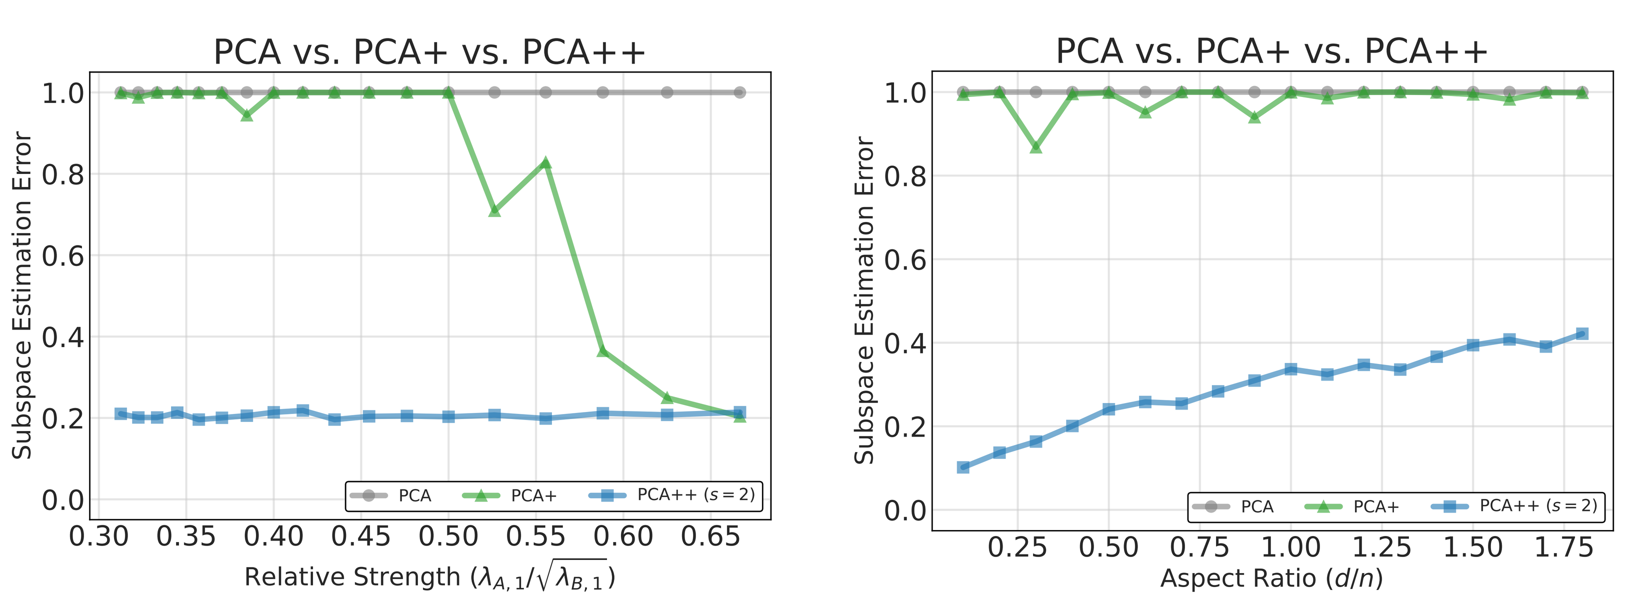

Evidence: PCA++, remains robust under strong background noise.

4. Our Method: PCA++ (Alignment + Uniformity)

Our method is a linear instantiation of the general spectral contrastive learning objective $\mathcal{L}(f)$, which balances alignment and uniformity terms:

$$\mathcal{L}(f) = \underbrace{-\mathop{\mathbb{E}}_{(x,x^+)}\left[f(x)^\top f(x^+)\right]}_{\text{alignment}} + \tau \cdot \underbrace{\left\| \mathop{\mathbb{E}}_{x}\left[f(x)f(x)^\top\right] - I_k \right\|_F}_{\text{uniformity}}. \tag{1}$$

Alignment + Uniformity = Robust Signal Recovery

Our PCA++ Formulation: For a linear encoder $f(x) = \mathbf{V}^\top x$ and enforcing a hard uniformity constraint ($\tau \to \infty$):

$$\max_{V \in \mathcal{R}^{n\times k}} \underbrace{\mathrm{tr}(V^\top S_n^+ V)}_{\text{alignment}}, \qquad \text{s.t.} \qquad \underbrace{V^\top S_n V = I_k}_{\text{uniformity}}, \tag{2}$$

Where $S_n^+ = \frac{1}{2n}(X^\top X^+ + X^{+\top}X)$ is the contrastive covariance, S_n is the standard sample covariance, and k is the number of signal dimensions.

This is solved efficiently via a **generalized eigenvalue problem**.

5. Theoretical Failure of PCA+ (Alignment-Only) Methods

Theorem 3.2 (Error Bound Deterioration): The subspace recovery error of PCA+ is explicitly bounded by:

$$\mathrm{dist}(\hat{U}_A, U_A) \le \frac{1}{\lambda_{A,k}}\left(\dots + \sqrt{\lambda_{A,1}\lambda_{B,1}\frac{m}{n}} + \dots\right).$$

Implication: As background noise $\lambda_{B,1}$ increases, the error is guaranteed to grow, leading to inevitable failure.

Theorem 3.4 (Catastrophic Misalignment): In a one-signal, one-background setting, we prove the alignment of the PCA+ estimator $\hat{v}_1$ with the true signal e_1:

$$\lim_{n,d\to\infty} \langle \hat{v}_1, e_1 \rangle^2 \le \frac{\lambda_{A,1}}{\sqrt{\lambda_{B,1}p/n}}.$$

Implication: When the background is strong enough ($\lambda_{B,1}$ is large), the right-hand side can become zero, proving that the estimator **fails to recover the signal.**

Conclusion: These theorems establish that a uniformity-type constraint is **theoretically necessary** for robust signal recovery under strong structured noise.

6. Theoretical Guarantees for PCA++

Theorem 4.2 (Fixed Aspect Ratio Regime)
As $n, d \to \infty$ with $d/n \to c > 0$, the squared subspace distance to the true signal subspace U_A converges to a closed-form limit:

$$\mathrm{dist}^2(\hat{U}_A, U_A) \to 1 - \frac{1 - c\lambda_{A,k}^{-2}}{1 + c\lambda_{A,k}^{-2}}.$$

Implication: PCA++ successfully recovers the signal even when background noise is arbitrarily strong, with a predictable, controlled error.

Theorem 4.4 (Growing Spike Regime)
When signal spikes grow proportionally to the dimension ($n\lambda_{A,k}/d \to c_{A,k}$), the limiting error statisfies:

$$\mathrm{dist}^2(\hat{U}_A, U_A) \to \frac{c_{A,k}}{1 + c_{A,k}}.$$

Implication: The error depends only on the effective SNR, confirming the method's robustness and consistency across different asymptotic regimes.

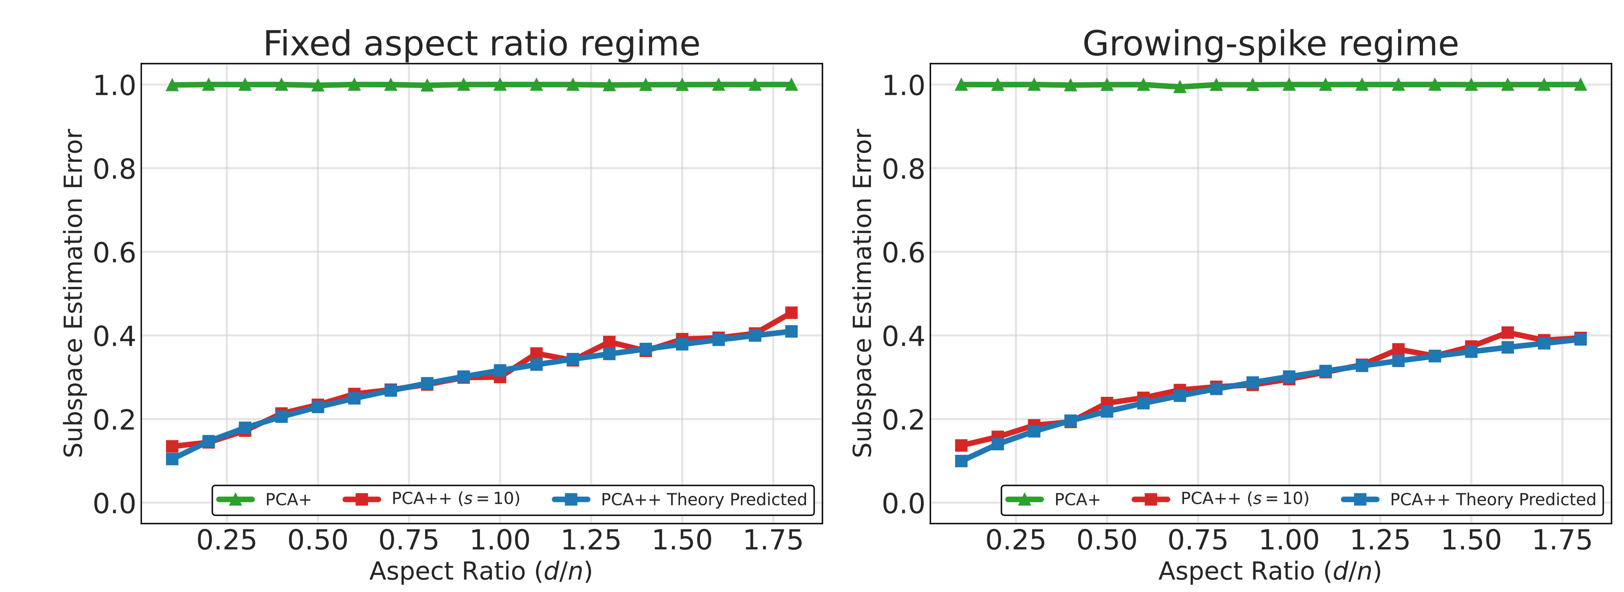

Theory Matches Practice: In both regimes, our theoretical predictions (red) perfectly align with experimental results (blue).

7. Key Takeaways

- **Alignment alone is insufficient** for robust signal recovery.
- **Uniformity is a key mechanism for robustness** against structured noise.
- **PCA++ is a practical and provably robust method** where others fail.
- Our work clarifies the **statistical role** of uniformity in contrastive learning.