Eigenvector Phase Transitions under Anisotropic Noise

Anonymous authors
Paper under double-blind review

Abstract

Identifying latent structures in environmental data—such as habitat clusters or pollution sources—is a fundamental challenge in ecological and climate science. Spectral methods, which analyse the principal eigenvectors of affinity matrices, are powerful tools for this task. However, environmental systems are rarely isotropic; physical processes like river flows or prevailing winds create strong directional gradients, resulting in anisotropic noise. The effect of such anisotropy on the reliability of spectral methods is not yet well understood in the literature. In this work, we develop a rigorous theory for this scenario by analysing a spiked random matrix model subjected to anisotropic noise. We derive an exact, analytical expression for the critical signal-to-noise ratio required for signal detection, establishing a sharp phase transition. Our central result proves that this threshold depends critically on the geometric alignment between the signal and the dominant environmental gradient, formalising a "camouflage effect". We also uncover a critical failure mode where this environmental gradient can itself create a "phantom" structure that spectral methods can easily detect, posing a significant risk of misinterpretation for scientists. Furthermore, we show that in the detectable phase, the second eigenvector aligns with the primary noise direction, revealing a deeper reorganisation of the system's structure. We complete our analysis with a Central Limit Theorem for the alignment fluctuations. We validate our theoretical predictions with simulations of ecological systems, offering a fundamental understanding of when spectral methods succeed or fail in realistic environments. Code to reproduce all results in the paper is anonymously released at https://anonymous.4open.science/r/tmlr_ept

1 Introduction

Spectral methods, which leverage the eigenvectors of affinity or covariance matrices, are a cornerstone of modern data analysis, with profound applications in machine learning and natural sciences (von Luxburg, 2007; Ng et al., 2001). In environmental science, these techniques are indispensable for uncovering latent structures from complex datasets, such as identifying distinct ecological communities from species abundance data, delineating habitat corridors from genetic information, or isolating dominant modes of climate variability (Legendre & Legendre, 2012). The success of these methods hinges on a critical assumption: *that the principal eigenvectors of a data-derived matrix faithfully align with the underlying structure of interest.* However, the reliability of this alignment can be severely compromised by noise, and the nature of this noise in environmental systems is often far from simple.

The foundational theory for understanding the limits of signal detection in noise comes from random matrix theory, particularly the study of spiked matrix models (Baik et al., 2005; Johnstone, 2001). The seminal work on the BBP phase transition revealed that for a low-rank signal matrix perturbed by uniform, isotropic noise, a sharp signal-to-noise threshold exists below which the signal is statistically undetectable (Baik et al., 2005; Johnstone, 2001). This theory provides a fundamental understanding of when spectral methods fail, showing that even an arbitrarily strong signal can be lost if the matrix size is large enough. While these results are powerful, they are built on the assumption of isotropic noise, where random fluctuations are equally likely in all directions. This assumption is frequently violated in environmental contexts. Physical processes such as prevailing winds, river currents, or geological formations impose a strong directional "grain" on the system, leading to anisotropic noise. The critical gap in the literature is a rigorous, analytical theory

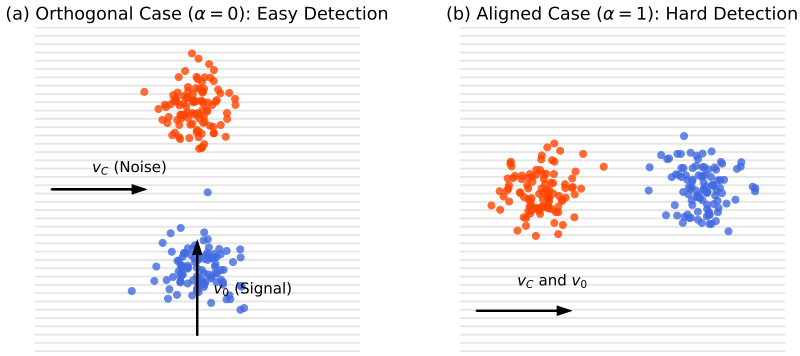

Figure 1: A conceptual diagram of the "camouflage effect" investigated in this paper. In both panels, the background texture represents a dominant environmental gradient (noise direction v_C). The signal consists of two latent communities (blue and red points). **(a)** When the signal structure (v_0) is orthogonal to the noise, the communities are easily distinguishable. In this regime, the signal-to-noise ratio required for detection is low, as defined by the threshold $\beta_c(0) = c_{\text{weak}}$. **(b)** When the signal is aligned with the noise, it is "camouflaged" by the environmental gradient. This makes the structure much harder to detect and requires a significantly higher signal strength, as defined by the threshold $\beta_c(1) = c_{\text{strong}}$.

that describes how such structured, anisotropic noise affects the performance of spectral methods. It is not known, for instance, how the geometric alignment between a latent signal and a dominant environmental gradient influences the threshold for detectability.

To address this gap, we analyse a spiked random matrix model where the noise is explicitly anisotropic. Our model consists of a rank-one signal matrix, representing a latent two-community structure, perturbed by a random noise matrix whose covariance is non-uniform, possessing a single dominant direction. This setup is mathematically tractable yet environmentally justified, directly modelling scenarios where a primary physical process creates a directional bias (Anderson et al., 2012). The alignment between the signal and the dominant noise direction is captured by a single geometric parameter, α, as conceptually illustrated in Figure 1.

Within this framework, we develop a complete theoretical characterisation of the system's behaviour. We prove the existence of a sharp phase transition for signal detection and derive an exact, analytical formula for the critical threshold. Our central finding is that this threshold is a non-trivial function of the alignment α, proving that a signal is significantly harder to detect when it is aligned with the environmental grain—a phenomenon we term the "camouflage effect". We extend this analysis to show that above the detection threshold, the system's eigenvectors undergo a remarkable reorganisation: the principal eigenvector aligns with the signal, while the second eigenvector aligns with the dominant noise direction. This reveals that spectral methods can, in principle, disentangle *both* the latent structure and the primary source of environmental noise. Finally, we establish a Central Limit Theorem for the fluctuations of the eigenvector alignment, providing a precise measure of the detection uncertainty, which we show is maximised at the critical point.

Our theoretical claims are validated through a comprehensive suite of numerical experiments. We first use direct simulations of our theoretical model to verify each of our theorems with high precision. We then conduct an experiment using a more realistic, non-linear simulation of an ecological system to demonstrate that the insights from our idealised model are robust and hold in a setting analogous to real-world applications.

The key contributions of this paper are:

- A rigorous, analytical formula for the phase transition threshold for signal detection in the presence of anisotropic noise, which explicitly depends on the signal-noise alignment.
- The discovery and formalisation of an eigenvector reorganisation phenomenon, where the second eigenvector captures the dominant noise direction in the detectable phase.

- A Central Limit Theorem for the eigenvector alignment fluctuations, providing a precise characterisation of the detection uncertainty.

- A comprehensive experimental validation that confirms our theoretical predictions and demonstrates their relevance to scientific discovery in environmental science.

The remainder of this paper is organised as follows. In §2, we formally define our mathematical model. In §3, we present the derivation of our main theoretical results. The experimental validation plan and its results are detailed in §4. We conclude with a discussion of related work and the implications of our findings in §5.

Table 1: Key symbols and notations in the paper.

Symbol	Description
Model Parameters	
N	Dimension of the matrix (number of nodes/points)
A	The full $N \times N$ observed matrix ($A = A_{\text{signal}} + W_{\text{aniso}}$)
W	Baseline $N \times N$ Wigner matrix (i.i.d. noise)
C	Deterministic $N \times N$ noise covariance structure matrix
W_{aniso}	The final $N \times N$ anisotropic noise matrix
v_0	Unit vector representing the ground-truth signal direction
v_C	Principal eigenvector of C, representing the dominant noise direction
β	Scalar signal strength ($\beta > 0$)
α	Alignment parameter, $\alpha = \lvert\langle v_0, v_C\rangle\rvert \in [0, 1]$
c_{strong}	Principal eigenvalue of C (strength of directional noise)
c_{weak}	Eigenvalue of C for background noise components
Theoretical Quantities	
τ	Spectral edge of the noise bulk ($\tau = 2c_{\text{weak}}$)
$\beta_c(\alpha)$	Critical threshold for signal detection
$\hat{\lambda}_k$	The k-th largest eigenvalue of the observed matrix A
$\hat{v}_k$	The eigenvector corresponding to $\hat{\lambda}_k$
$f(\beta, \alpha)$	Asymptotic (mean) alignment $\lvert\langle \hat{v}_1, v_0\rangle\rvert^2$ for $\beta > \beta_c$
$\sigma^2(\beta, \alpha)$	Asymptotic variance of the alignment fluctuations

2 The Model

In this section, we formally define the anisotropic spiked random matrix model used for our theoretical analysis. The model is constructed to capture the essential features of a latent signal embedded within a noisy environment that possesses a dominant directional structure. For clarity, the key notations introduced in this paper are summarised in Table 1.

2.1 The Signal Matrix

The signal component of our model is a deterministic, rank-one matrix representing the ground-truth latent structure we aim to recover. It is defined as:

$$A_{\text{signal}} = \beta v_0 v_0^T \tag{1}$$

where $\beta > 0$ is a non-random scalar representing the signal strength. The vector $v_0 \in \mathbb{R}^N$ is a deterministic unit vector ($\|v_0\|_2 = 1$) that encodes the structure. For the canonical problem of detecting a two-community structure in a network of N nodes, this vector takes the form $v_0 = \frac{1}{\sqrt{N}}[1, \ldots, 1, -1, \ldots, -1]^T$, where the two communities are of equal size. The principal eigenvector of A_{signal} is precisely v_0, and our central goal is to determine under what conditions the principal eigenvector of the full, noisy matrix aligns with it.

2.2 The Anisotropic Noise Matrix

The noise component is designed to model random fluctuations that are not uniform, but are instead shaped by an underlying environmental structure. Its construction begins with a baseline of uniform noise, which is then deformed by a deterministic covariance structure.

Baseline Noise. We begin with a standard $N \times N$ Wigner matrix, W. This is a symmetric random matrix where the entries W_{ij} for $i \leq j$ are independent and identically distributed random variables with mean $\mathbb{E}[W_{ij}] = 0$ and variance $\mathbb{E}[W_{ij}^2] = 1/N$. This normalisation ensures that the eigenvalues of W converge to the well-known Wigner semicircle law, supported on the interval $[-2, 2]$ as $N \to \infty$ (Anderson et al., 2009).

Covariance Structure. The anisotropy is introduced via a deterministic, positive semi-definite matrix $C \in \mathbb{R}^{N \times N}$, which represents the covariance profile of the noise. To maintain tractability while capturing the essence of anisotropy, we assume a simple yet powerful structure for C. We assume that it has one large eigenvalue, c_{strong}, corresponding to a principal eigenvector v_C, and that its remaining $N-1$ eigenvalues are identical, equal to a background value c_{weak}. This structure is a suitable model for systems with a single dominant environmental gradient (Soons et al., 2004; Hughes et al., 2009).

Final Construction. We construct the final anisotropic noise matrix, W_{aniso}, by deforming the baseline Wigner matrix with the covariance structure C:

$$W_{\text{aniso}} = C^{1/2} W C^{1/2} \tag{2}$$

This construction ensures that the statistical properties of the noise are shaped by C, with greater variance in the directions corresponding to larger eigenvalues of C. This method of modelling structured noise is a standard approach in the study of deformed random matrix ensembles (Benaych-Georges & Nadakuditi, 2011).

2.3 The Full Model and Alignment Parameter

Combining the signal and noise components gives us the final observable matrix A:

$$A = \beta v_0 v_0^T + W_{\text{aniso}} \tag{3}$$

The crucial interaction between the signal's structure and the noise's structure is captured by a single geometric parameter. We define the alignment parameter, α, as the absolute inner product of the signal direction and the dominant noise direction:

$$\alpha = |\langle v_0, v_C \rangle| \tag{4}$$

This parameter measures the collinearity of the two directions, ranging from $\alpha = 0$ when the signal is perfectly orthogonal to the environmental grain (Figure 1 (a)), to $\alpha = 1$ when it is perfectly aligned (Figure 1 (b)).

3 Theoretical Analysis

In this section, we present the analytical core of our work. We begin by characterising the eigenvalue spectrum of the anisotropic noise matrix, which establishes the baseline for our analysis (§3.1). We then analyse the full signal-plus-noise model to derive the sharp phase transition for signal detection and quantify the quality of the eigenvector alignment (§3.2). Building on this, we reveal a deeper reorganisation of the eigenspace involving the second eigenvector (§3.3). Finally, we establish a Central Limit Theorem for the fluctuations of the eigenvector alignment, providing a precise measure of detection uncertainty (§3.4). The detailed proofs of these theoretical claims are given in §A.

3.1 The Noise Spectrum and Spectral Edge

To understand the conditions under which a signal can be detected, we must first characterise the eigenvalues of the noise matrix W_{aniso} in isolation. These eigenvalues form a continuous "bulk" from which a signal-aligned eigenvalue must emerge to be detectable. Our first theorem defines the boundary of this bulk.

Theorem 1 (The Spectral Edge)**.** *In the limit $N \to \infty$, the continuous part of the eigenvalue spectrum of the anisotropic noise matrix W_{aniso} is supported on the compact interval $[-\tau, \tau]$, where the spectral edge τ is given by:*

$$\tau = 2c_{weak} \tag{5}$$

Proof Sketch. The proof relies on the stability of the continuous spectrum of a large random matrix under a finite-rank perturbation. We decompose our covariance structure matrix as $C = c_{\text{weak}}I + (c_{\text{strong}} - c_{\text{weak}})v_C v_C^T$. The noise matrix W_{aniso} can then be seen as a base isotropic noise matrix, $c_{\text{weak}}W$, perturbed by a term related to the rank-one spike of C. A key result in random matrix theory states that such a perturbation can pull out discrete outlier eigenvalues but does not change the edges of the continuous spectrum (Benaych-Georges & Nadakuditi, 2011). The spectrum of $c_{\text{weak}}W$ is a scaled Wigner semicircle supported on $[-2c_{\text{weak}}, 2c_{\text{weak}}]$, which therefore defines the edges of the bulk for W_{aniso}. See Appendix A.2 for the detailed proof.

Interpretation. This result is powerful and non-trivial. It means that the primary barrier to detecting a new signal is determined *not* by the strongest, most obvious component of the environmental noise (c_{strong}), but by the magnitude of the uniform, background fluctuations (c_{weak}). The dominant anisotropic noise may create its own outlier eigenvalue, but the "sea" of noise that a new signal must rise above is defined by the weaker, isotropic component.

3.2 The Phase Transition for Signal Detection

A signal is statistically detectable via spectral methods if and only if its corresponding eigenvalue, λ_1, "pops out" from the noise bulk, i.e., $\lambda_1 > \tau$. We find that this occurs only when the signal strength β exceeds a critical threshold, which depends fundamentally on the alignment α.

Theorem 2 (The Critical Threshold)**.** *An isolated eigenvalue corresponding to the signal v_0 emerges from the noise bulk if and only if the signal strength β exceeds a critical threshold $\beta_c(\alpha)$, given by:*

$$\beta_c(\alpha) = \frac{1}{\frac{\alpha^2}{c_{strong}} + \frac{1-\alpha^2}{c_{weak}}} \tag{6}$$

Proof Sketch. We analyse the characteristic equation for outlier eigenvalues derived from the resolvent method (see Lemma 1 in Appendix A.1). The phase transition occurs at the minimum value of β for which a solution $\lambda_1 > \tau$ exists. By setting $\lambda_1 = \tau = 2c_{\text{weak}}$ and evaluating the large-N limit of the resolvent term $\langle v_0, (W_{\text{aniso}} - \tau I)^{-1} v_0 \rangle$, we can solve for this minimum β. The calculation involves decomposing v_0 into components parallel and orthogonal to the noise direction v_C, which naturally introduces the dependence on α. The detailed proof is provided in Appendix A.3.

Interpretation. This provides a rigorous quantification of the "camouflage effect", effectively translating the standard concept of a direction-dependent signal-to-noise ratio into the framework of random matrix theory. The denominator represents the effective noise power in the direction of the signal. When the signal is aligned with the environmental grain ($\alpha = 1$), it must overcome the strong noise component ($\beta_c = c_{\text{strong}}$). When it is orthogonal ($\alpha = 0$), it only needs to overcome the weak background noise ($\beta_c = c_{\text{weak}}$). It is therefore harder to detect a latent structure that aligns with a dominant environmental gradient.

Once the signal is detectable, we can quantify the quality of the recovery.

Theorem 3 (Asymptotic Alignment)**.** *For a signal strength $\beta > \beta_c(\alpha)$, the squared inner product (alignment) between the principal eigenvector of A, $\hat{v}_1$, and the true signal vector v_0 converges to:*

$$|\langle \hat{v}_1, v_0 \rangle|^2 = 1 - \frac{\beta_c(\alpha)^2}{\beta^2} \tag{7}$$

Proof Sketch. The alignment is derived by analysing the residue of the resolvent at the pole λ_1. The magnitude of the eigenvector's projection onto v_0 is functionally related to the location of the eigenvalue λ_1, which is itself a function of β. Solving this system of equations for the projection magnitude yields the stated result. The detailed proof is provided in Appendix A.3.

Interpretation. This result shows that alignment is zero at the exact moment of the phase transition and increases towards perfect recovery ($|\langle \hat{v}_1, v_0 \rangle|^2 \to 1$) as the signal strength β grows.

3.3 Eigenspace Reorganisation: The Second Eigenvector

The introduction of a strong signal does more than just create a new top eigenvector; it reorganises the entire eigenspace. The next most significant feature of the system—the dominant environmental gradient—is displaced from the first to the second eigenvector.

Theorem 4 (Second Eigenvector Alignment)**.** *In the super-critical phase ($\beta > \beta_c(\alpha)$), the second eigenvector of A, $\hat{v}_2$, aligns with the principal noise direction v_C.*

Proof Sketch. The proof involves analysing the characteristic equation for outlier eigenvalues when the noise matrix W_{aniso} is itself considered a spiked matrix (with a spike corresponding to v_C). The interaction between the signal spike and the noise spike creates two outlier eigenvalues. By analysing the structure of the eigenvectors corresponding to these two solutions, we can show that the larger eigenvalue corresponds to an eigenvector aligned with v_0, while the smaller of the two corresponds to an eigenvector aligned with v_C. The detailed proof is given in Appendix A.4.

Interpretation. This spectral "sorting" phenomenon is a subtle and powerful result. It implies that a careful spectral analysis can simultaneously reveal both the primary latent signal and the primary axis of environmental noise. The method not only detects the structure but also provides information about the nature of the unkown interference, which is of scientific merit as well.

3.4 Fluctuations and Detection Uncertainty

Our final result moves beyond asymptotic averages to characterise the random fluctuations of the alignment for finite N, providing a measure of statistical confidence.

Theorem 5 (Central Limit Theorem for Alignment)**.** *In the super-critical phase ($\beta > \beta_c(\alpha)$), the fluctuations of the eigenvector alignment are asymptotically Gaussian. The scaled quantity converges in distribution to:*

$$\sqrt{N}\left(|\langle \hat{v}_1, v_0 \rangle|^2 - f(\beta, \alpha)\right) \xrightarrow{d} \mathcal{N}(0, \sigma^2(\beta, \alpha)) \tag{8}$$

where $f(\beta, \alpha)$ is the asymptotic alignment from Theorem 3, and $\sigma^2(\beta, \alpha)$ is a deterministic variance function given in Equation (66).

Proof Sketch. The proof is based on establishing a CLT for the fluctuations of the resolvent term $\langle v_0, (W_{\text{aniso}} - \lambda_1 I)^{-1} v_0 \rangle$. The variance of this term can be computed using established methods from random matrix theory. This variance is then propagated through a linearised model of the system's characteristic equations to find the resulting variance of the eigenvector alignment. The proof relies on showing that the fluctuations of the eigenvalue and eigenvector are linearly coupled. See Appendix A.5 for the detailed proof.

Interpretation. The variance function $\sigma^2(\beta, \alpha)$ provides a rigorous quantification of the reliability or statistical confidence of the signal detection, as conceptually illustrated in Figure 2. It measures the expected "jitter" in the alignment, telling an environmental scientist how stable the detected pattern is likely to be. A detailed analysis of its mathematical form, which is a function of the derivatives of the limiting resolvent $F(z)$, reveals several profound insights:

- **Instability at the Tipping Point:** The variance is maximised as the signal strength β approaches the critical threshold β_c. This is the mathematical signature of a critical phenomenon. As the outlier

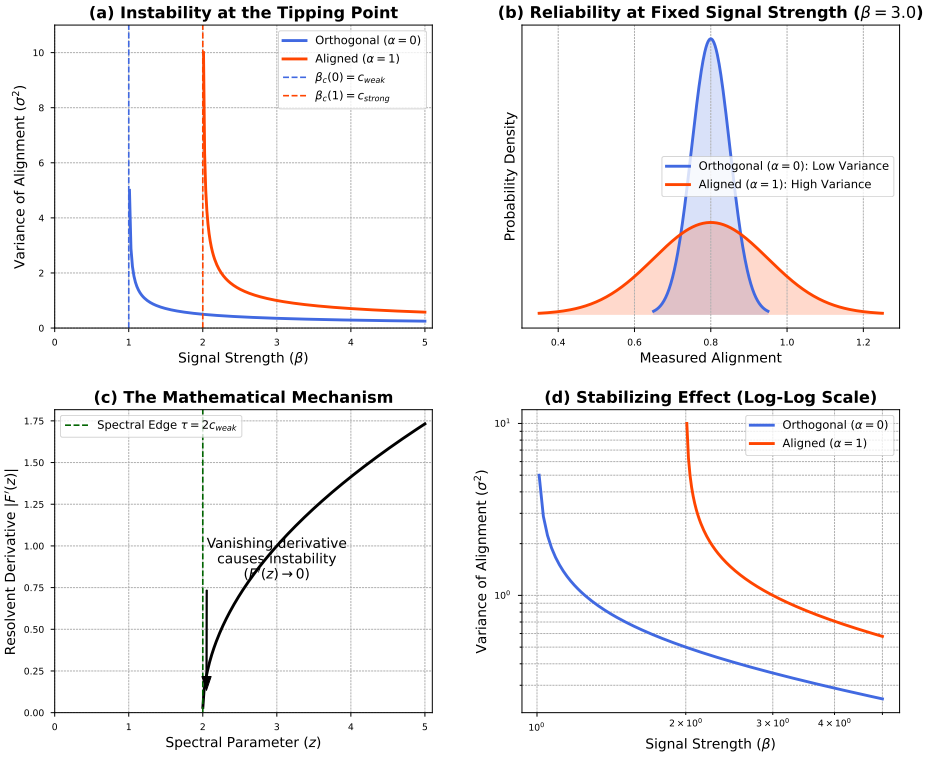

Figure 2: Illustration of the theoretical detection uncertainty, as characterised by the variance of the alignment fluctuations (σ^2). **(a)** The variance is plotted against signal strength (β), revealing a sharp divergence at the critical thresholds. This peak demonstrates the extreme instability of the system at its tipping point. **(b)** At a fixed super-critical signal strength ($\beta = 3.0$), the probability distribution of the measured alignment is shown. The measurement is more reliable (lower variance, taller peak) for the orthogonal case ($\alpha = 0$) than for the aligned case ($\alpha = 1$), highlighting the practical impact of the camouflage effect on detection confidence. **(c)** The underlying mathematical mechanism for the instability is revealed. The first derivative of the resolvent, $|F'(z)|$, is plotted against the spectral parameter z. The derivative vanishes precisely at the spectral edge (τ), causing the mathematical singularity that drives the divergence in variance. **(d)** The same data as in panel (a) is shown on a log-log scale. The near-straight-line decay for large β confirms that the system predictably stabilises as the signal becomes stronger, and the detection becomes more reliable.

eigenvalue $\bar{\lambda}_1$ approaches the spectral edge τ, the system becomes "soft", which is mathematically reflected by the first derivative of the resolvent function, $F'(\bar{\lambda}_1)$, approaching zero. This vanishing derivative means that the system has no "restoring force" against perturbations. Since the variance formula contains a high power of $F'(\bar{\lambda}_1)$ in its denominator, this mathematical singularity causes the variance to diverge, precisely capturing the physical intuition of extreme instability at a tipping point.

- **The Stabilising Effect of a Strong Signal:** As β grows much larger than the threshold, the variance decreases towards zero. For large β, the eigenvalue $\bar{\lambda}_1$ is pushed far from the noise bulk, into a region where the resolvent and its derivatives are smooth and well-behaved. Here, $F'(\bar{\lambda}_1)$ is non-zero and stable, meaning that the system is "stiff" and has a strong restoring force that suppresses fluctuations. The variance is then dominated by the $1/\beta^4$ term, rigorously showing that a strong signal stabilises the system and makes its detection highly reliable.

- **The Role of Alignment in Detection Reliability:** The alignment α influences the variance through the "driving noise" term, $\mathcal{V}(\bar{\lambda}_1)$. The calculation of this term requires decomposing the signal vector v_0 into components parallel and orthogonal to the dominant noise direction v_C. The

final variance is thus a weighted sum of the fluctuations originating from the strong and weak noise directions. This provides an justification for the intuition that the reliability of a measurement depends on its geometric relationship with the primary environmental gradients.

Ultimately, this theorem provides a powerful new tool. It allows scientists to assess the confidence in a detected pattern based not only on its apparent strength but also on its geometric relationship with the primary gradients of the environment itself.

4 Experimental Validation

In this section, we present a comprehensive suite of numerical simulations designed to validate the theoretical claims from §3. Our experimental validation is structured in two parts. First, we conduct a series of direct validations of our idealised signal model, providing a rigorous, one-to-one proof for each of our theorems. Second, to demonstrate that the insights from our linear model are not confined to that idealised setting, we test their applicability in an illustrative, non-linear simulation involving spectral clustering. Unless otherwise specified, the main results are generated using a matrix size of $N = 1000$, with 50 Monte Carlo trials per data point, and noise parameters of $c_{\text{strong}} = 10.0$ and $c_{\text{weak}} = 1.0$. This high 10:1 anisotropy ratio was chosen deliberately to create a strong and unambiguous experimental setting, allowing for a clear validation of the theoretical phenomena of interest, such as the "camouflage effect" and "noise hijacking".

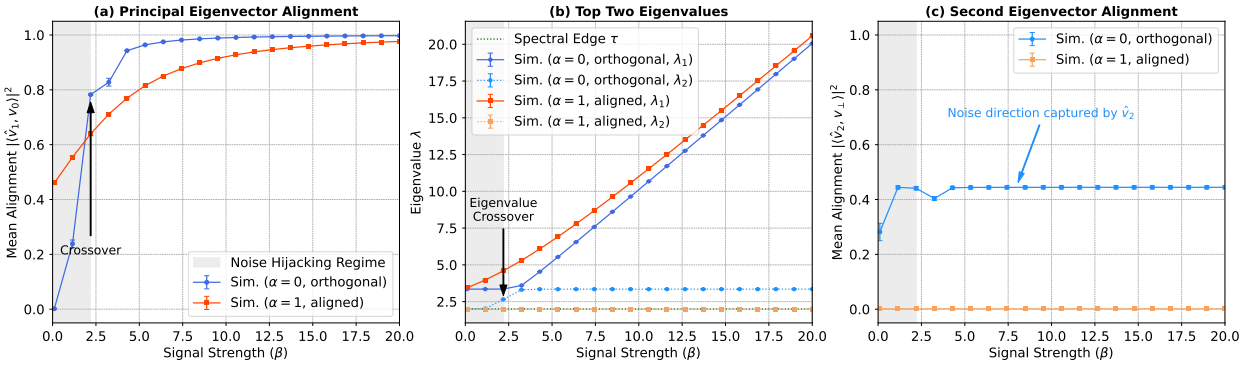

Figure 3: Validation of eigenvector phase transitions with finite-size Monte Carlo simulations ($N = 1000$, $c_{\text{strong}} = 10$, $c_{\text{weak}} = 1$, 50 trials per signal strength β for mean and error bar estimation). **(a) Principal Eigenvector Alignment:** This panel shows that for low signal strength, the alignment for the aligned case ($\alpha = 1$, red) is higher due to a "noise hijacking" effect. The orthogonal case ($\alpha = 0$, blue) undergoes its sharp phase transition at $\beta \approx 1$ and eventually crosses over the aligned case and approaches its theoretical limit of perfect signal detection. **(b) Top Two Eigenvalues:** This panel reveals that for the aligned case ($\alpha = 1$), the signal and noise merge into a single dominant eigenvalue (λ_1, solid red), which grows continuously with β. The second eigenvalue (λ_2, dashed red) remains flat at the spectral edge, $\tau = 2c_{\text{weak}}$. For the orthogonal case ($\alpha = 0$), the system supports two distinct outlier eigenvalues. Initially, the largest eigenvalue (λ_1, solid blue) corresponds to the noise spike, while the second eigenvalue (λ_2, dashed blue) corresponds to the rising signal. After a point of eigenspace reorganisation around $\beta \approx 2.5$, the signal becomes dominant, taking over the role of λ_1, while the noise settles into the role of λ_2. The gap between the two eigenvalues narrows near this reorganisation point before widening again. **(c) Second Eigenvector Alignment:** This panel shows the consequence of the eigenvalue behaviour. For the aligned case, the alignment of $\hat{v}_2$ is near zero, as there is no distinct second spike to capture. For the orthogonal case, the alignment of $\hat{v}_2$ with the noise direction rises during the unstable reorganisation phase and then stabilises at a high value once the signal has cleanly separated and taken over the principal eigenvector.

4.1 Direct Validation of the Theoretical Model

Our first goal is to provide numerical proof for our theoretical framework using simulations that exactly match the model defined in §2. The most important test is to validate the central predictions for the phase transition and the reorganisation of the eigenspace. The results, shown in Figure 3, provide a complete picture of the system's behaviour. Panel (a) confirms the existence of the "noise hijacking" effect, where the alignment for the aligned case ($\alpha = 1$) is high even for weak signals, and the subsequent crossover phenomenon. Panel (b) reveals the underlying mechanism: an "eigenvalue crossover" (an identity swap) between the rising signal eigenvalue and the stable noise eigenvalue in the orthogonal case ($\alpha = 0$). Panel (c) provides direct visual proof for Theorem 4, showing that the second eigenvector successfully capturing the orthogonal noise direction after this crossover.

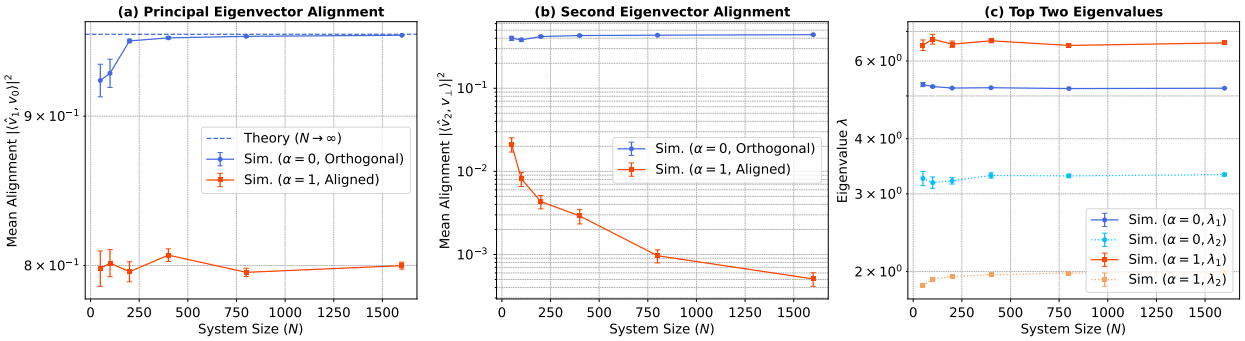

Figure 4: Validation of the asymptotic nature of eigenvector phase transitions and eigenspace reorganisation, showing the convergence of key observables (plotted in log scale) as the system size N increases ($N = [50, 100, 200, 400, 800, 1600]$, $c_{\text{strong}} = 10$, $c_{\text{weak}} = 1$, 50 trials per N for mean and error bar estimation). Both scenarios are run at a fixed signal strength of $\beta = 5.0$. **(a) Principal Eigenvector Alignment:** For the orthogonal case ($\alpha = 0$, solid blue line), the system is super-critical, and the alignment correctly converges to its theoretical limit (dashed blue line). For the aligned case ($\alpha = 1$), the system is sub-critical; the observed high alignment is a robust "noise hijacking" effect, where the eigenvector aligns with the dominant noise spike. **(b) Second Eigenvector Alignment:** For the orthogonal case, the alignment of $\hat{v}_2$ with the noise direction converges to a stable, non-zero value, confirming the asymptotic nature of the eigenspace reorganisation. For the aligned case, the alignment decays roughly as $\sim 1/N$, the expected behaviour for a random bulk eigenvector, confirming the merging of the signal and noise spikes. **(c) Top Two Eigenvalues:** The locations of all outlier eigenvalues stabilise rapidly as N increases, demonstrating their convergence to the predicted deterministic limits.

To rigorously validate our asymptotic theory, it is crucial to demonstrate that the finite-size effects observed in the previous figure diminish as the system size N increases. Figure 4 provides this proof. Panel (a) shows the principal eigenvector alignment from the simulation systematically converging towards the theoretical limit predicted by Theorem 3 as N grows. Panels (b) and (c) confirm that the more subtle features of the model—the second eigenvector alignment and the locations of the outlier eigenvalues—are also robust asymptotic properties that converge to their predicted states, solidifying the validity of our entire theoretical framework.

The final validation of our idealised model tests the predictions for the statistical fluctuations of the alignment. The results, shown in Figure 5 and Table 2, confirm the predictions of our Central Limit Theorem (Theorem 5). The histograms and Q-Q plots in Figure 5 provide clear visual evidence for the Gaussian nature of the fluctuations. The table then provides a quantitative analysis, revealing the deep physical insight that the variance is the highest for the intermediate case ($\alpha = 0.5$), as it operates closest to its critical tipping point, where system instability is maximised.

For completeness, we also validate the foundational assumption of our model: the structure of the anisotropic noise spectrum. Figure 6 confirms the predictions of Theorem 1. The figure illustrates the emergence of the

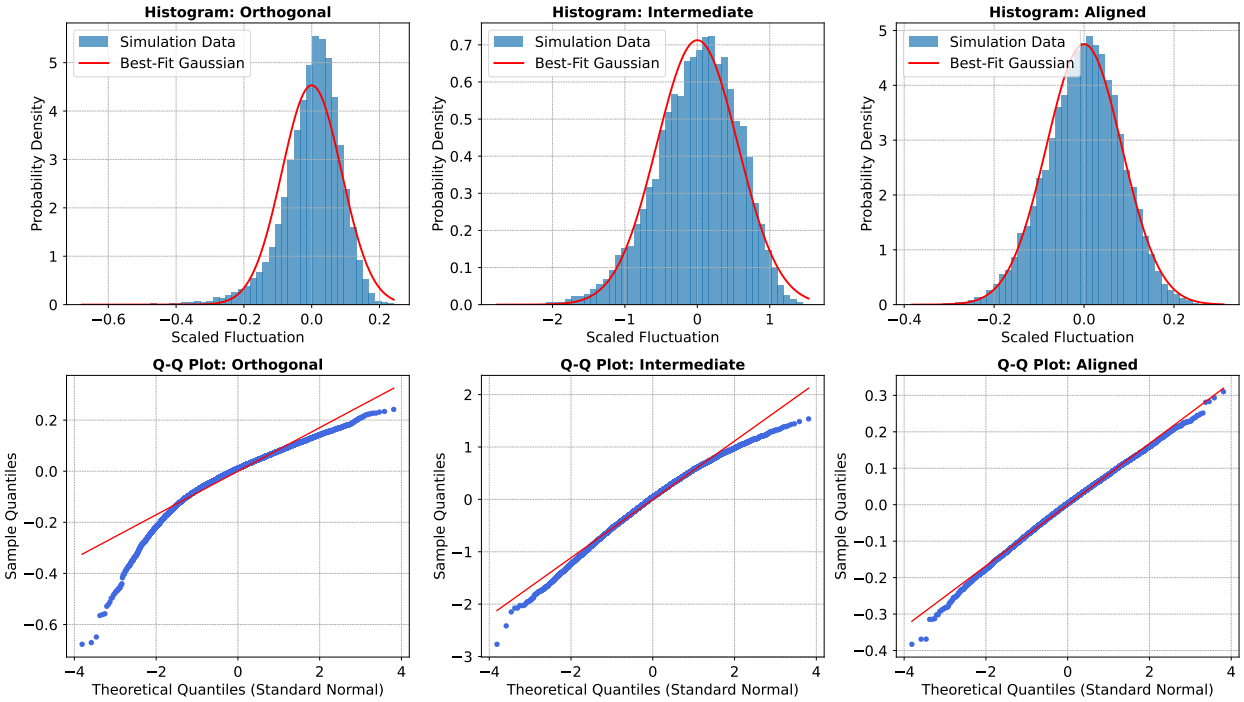

Figure 5: Validation of the Central Limit Theorem for alignment fluctuations at system size $N = 1000$ ($c_{\text{strong}} = 10$, $c_{\text{weak}} = 1$, $10,000$ Monte Carlo trials simulated). The top row shows histograms of the scaled fluctuations overlaid with a best-fit Gaussian, while the bottom row provides a more rigorous validation with Quantile-Quantile (Q-Q) plots. **Gaussianity:** In all three scenarios, the histograms closely match the Gaussian shape, and the Q-Q plots are nearly linear, providing strong evidence that the fluctuations are asymptotically Gaussian as predicted by Theorem 5. **Convergence Rate:** The quality of the Gaussian approximation reveals a deep and subtle insight into the model's physics. The convergence to the CLT is fastest for the aligned case ($\alpha = 1$), where the Q-Q plot is almost perfectly linear. This is due to the simpler physical system, where the signal and noise merge into a single spike. The convergence is slowest for the orthogonal case ($\alpha = 0$), where the Q-Q plot shows deviations in the tails. This reflects the more complex three-way interaction between the distinct signal spike, noise spike, and the random bulk, which generates stronger non-Gaussian finite-size effects that persist at $N = 1000$. The intermediate case ($\alpha = 0.5$) exhibits an intermediate convergence speed, consistent with this interpretation.

outlier "noise spike" as anisotropy increases, and the zoomed-in Panel (d) provides a high-precision validation that the continuous bulk of the spectrum terminates exactly at the theoretical edge, $\tau = 2c_{\text{weak}}$.

4.2 Application to Spectral Clustering in an Environmental Model

To bridge the gap between our idealised theory and a practical application, we now test our insights in a more realistic, non-linear setting. We simulate a 2D environmental model where two communities of points are generated, and we apply spectral clustering using an Radial Basis Function (RBF) kernel with an adaptive 'gamma' parameter to recover the latent structure.

Figure 7 provides a qualitative visualisation of the clustering performance. It clearly demonstrates the "camouflage effect" in a practical scenario: in the top row (aligned case, $\alpha = 1$), the algorithm fails to find the correct clusters even as signal strength increases, as all tested β values are below the high theoretical threshold of $\beta_c = 10.0$. In contrast, the bottom row (orthogonal case, $\alpha = 0$) shows a clear phase transition from failure to success as β crosses its low theoretical threshold of $\beta_c = 1.0$.

Table 2: Measured statistics of the scaled alignment fluctuations from the CLT experiment ($N = 1000$, $c_{\text{strong}} = 10$, $c_{\text{weak}} = 1$, $10,000$ trials). The table provides a quantitative validation of the Central Limit Theorem. **Mean:** In all scenarios, the mean of the centered fluctuations is correctly measured to be zero, as expected by construction. **Variance:** The variance quantifies the instability of the alignment measurement and reveals a deep physical insight. The variance is not a simple monotonic function of the alignment α. The intermediate case ($\alpha = 0.5$) exhibits a much larger variance (0.3134) than the extreme cases. This is because it is operating closer to its critical tipping point than the other two scenarios. For the intermediate case, the signal strength is $\beta = 5.0$ while its threshold is $\beta_c \approx 1.29$. For the orthogonal case, $\beta = 5.0$ is comparatively further from its threshold of $\beta_c = 1.0$. For the aligned case, $\beta = 15.0$ is also far from its threshold of $\beta_c = 10.0$. This result provides strong numerical evidence for the theoretical prediction that system instability is significantly higher near a critical phase transition.

Scenario	Measured Mean	Measured Variance
Orthogonal ($\alpha = 0.0$)	0.0000	0.0078
Intermediate ($\alpha = 0.5$)	0.0000	**0.3134**
Aligned ($\alpha = 1.0$)	0.0000	0.0071

Table 3: Comparison of the theoretical critical threshold from our idealised linear model with the empirically measured threshold from the non-linear spectral clustering experiment ($N = 1000$, $c_{\text{strong}} = 10$, $c_{\text{weak}} = 1$, 50 trials per signal strength to estimate mean). The measured threshold is defined as the signal strength β at which the mean Adjusted Rand Index (ARI) first exceeds 0.95. **Agreement:** For the orthogonal and intermediate cases, the measured thresholds are remarkably close to the theoretical predictions, demonstrating the predictive power of our theory even in a more complex, non-linear setting. **Discrepancy:** For the aligned case, the measured threshold is significantly lower than the theoretical prediction. This is a direct consequence of the "noise hijacking" effect; the spectral clustering algorithm achieves a high ARI score by correctly clustering the dominant pattern created by the strong anisotropic noise itself, long before the true signal is theoretically detectable. This highlights a crucial practical implication: strong, aligned environmental gradients can create detectable structures that may be misinterpreted as a true signal.

Scenario	Theoretical β_c	Measured β_c (ARI > 0.95)
Orthogonal ($\alpha = 0.0$)	1.00	1.10
Intermediate ($\alpha = 0.5$)	1.29	1.60
Aligned ($\alpha = 1.0$)	**10.00**	**2.50**

The quantitative performance of the clustering is shown in Figure 8, which plots the Adjusted Rand Index (ARI) against signal strength. This figure provides a rigorous, quantitative proof of the "camouflage effect". The performance curves for the three alignment scenarios ($\alpha = 0, 0.5, 1$) exhibit sharp phase transitions in the exact order predicted by our theory, with the aligned case requiring a much stronger signal to achieve successful clustering.

Finally, we connect the performance of this complex, non-linear application directly back to our simple, idealised linear theory. Table 3 compares the theoretical critical threshold, β_c, with the empirically measured threshold from the non-linear spectral clustering experiment. For the orthogonal and intermediate cases, there is a remarkable agreement between the theoretical prediction and the measured result. The table also highlights an important insight for the aligned case: the measured threshold is much lower than the theoretical one. This discrepancy is a direct result of the "noise hijacking" effect, where the algorithm achieves a high ARI score by correctly clustering the dominant pattern created by the strong anisotropic noise itself, long before the true signal is theoretically detectable. This demonstrates the powerful predictive utility of our theoretical framework in explaining the behaviour of real-world spectral methods.

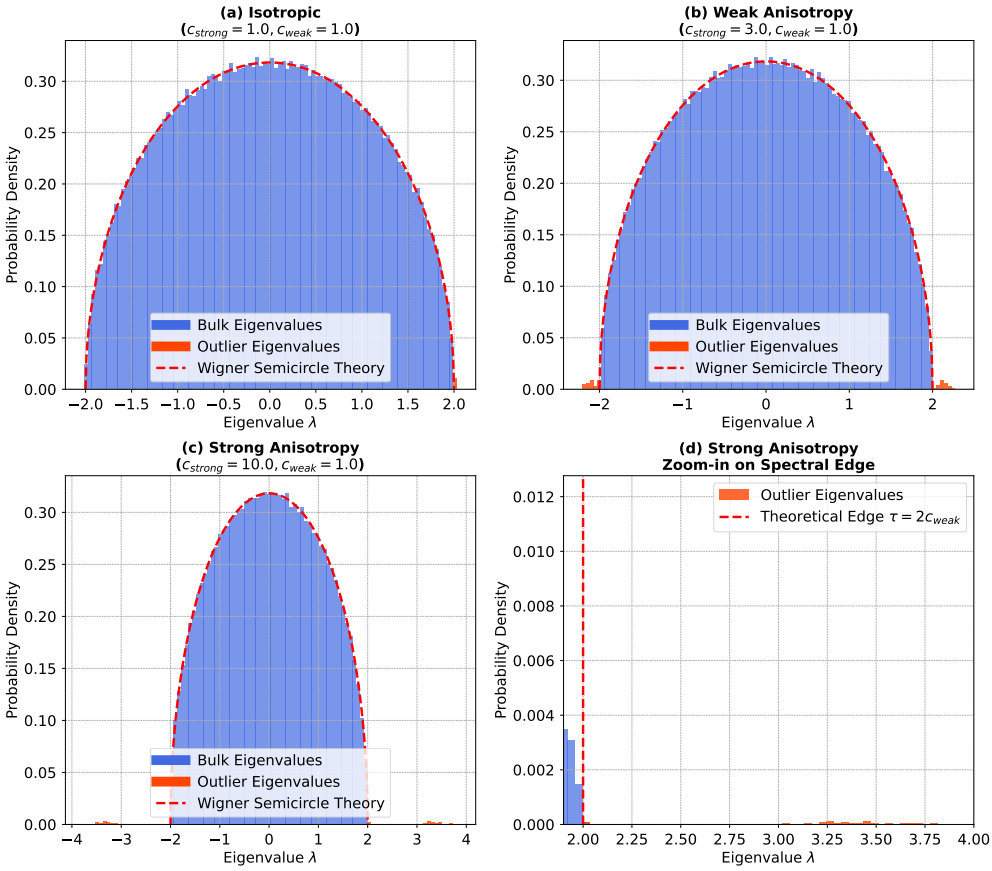

Figure 6: Validation of the anisotropic noise spectrum (Theorem 1) for $N = 1000$. The figure shows the empirical eigenvalue distribution for increasing levels of anisotropy. **(a)** In the isotropic case ($c_{\text{strong}} = c_{\text{weak}}$), the spectrum correctly matches the theoretical Wigner semicircle distribution. **(b)** With weak anisotropy, small but distinct outlier eigenvalues (the "noise spike", red) separate from the main bulk. **(c)** With strong anisotropy, the noise spike produces well-separated outliers, clearly distinct from the continuous bulk. **(d)** A zoomed-in view of the spectral edge for the strong anisotropy case provides a high-precision validation, showing that the continuous part of the spectrum (the bulk) terminates exactly at the theoretical prediction of $\tau = 2c_{\text{weak}}$ (red dashed line).

5 Discussion and Conclusion

In this work, we have developed and validated a complete theory for principal eigenvector phase transitions in the presence of anisotropic noise. Our findings extend the classical understanding of spectral methods by incorporating the crucial and realistic element of a dominant environmental gradient. We have shown that the detectability of a latent signal is not merely a function of its strength, but is fundamentally governed by its geometric alignment with the surrounding noise structure.

5.1 Implications of Key Findings for Scientific Discovery

Our theoretical and experimental results reveal several deep physical phenomena with direct practical implications. The central finding is the "camouflage effect" (Theorem 2), which proves that a latent structure is fundamentally harder to detect when it aligns with a strong environmental gradient, such as a river valley or prevailing wind. This has important consequences for experimental design and data interpretation in the environmental sciences, suggesting that the orientation of a study area relative to its dominant gradients can significantly impact the statistical power to detect underlying patterns.

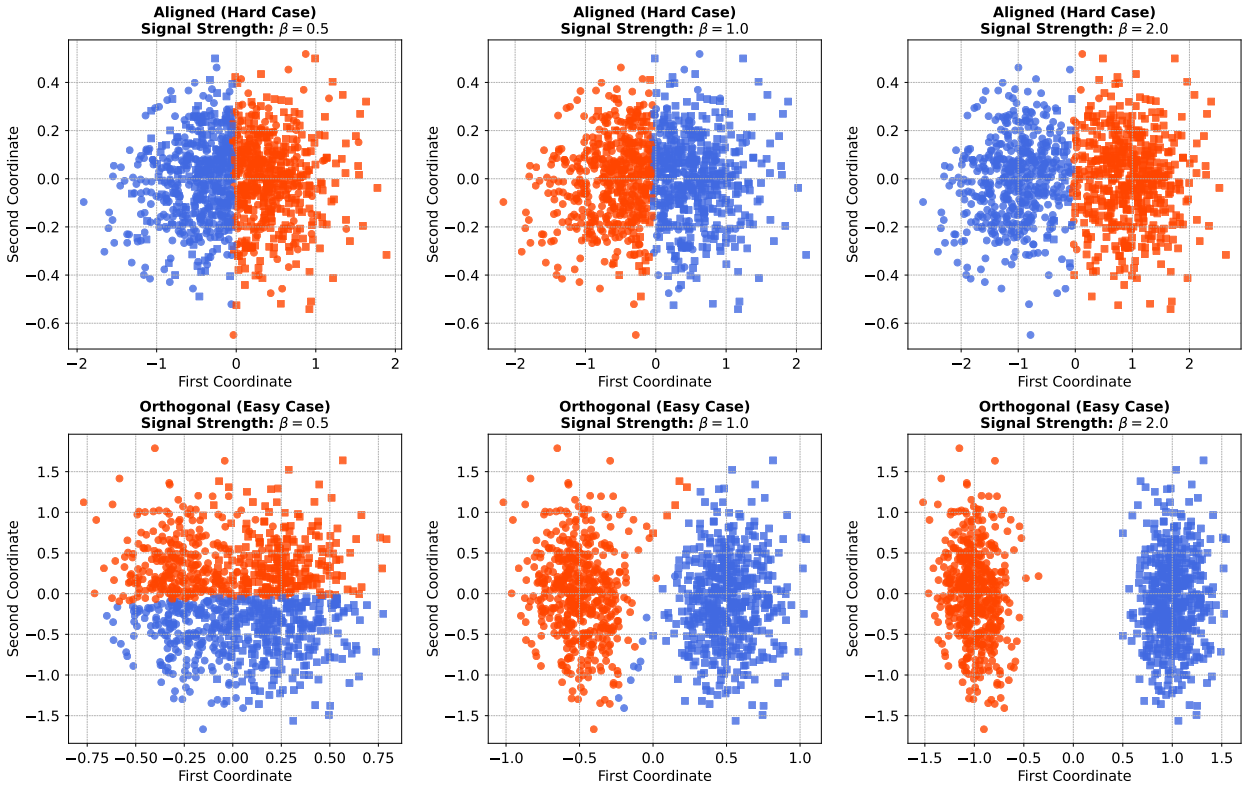

Figure 7: A visual demonstration of non-linear spectral clustering performance on the environmental model ($N = 1000$, $c_{\text{strong}} = 10$, $c_{\text{weak}} = 1$) across different signal strengths (β) and alignments (α). In each panel, the true community is indicated by the marker shape (circle or square), while the discovered community is indicated by the colour (blue or red). **Top Row (Aligned Case, $\alpha = 1$):** The clustering fails clearly in low signal strength regimes and succeeds only partially at $\beta = 2.0$. Even as the signal strength increases, the algorithm is unable to find the correct horizontal separation because the signal is "camouflaged" by the strong, aligned environmental noise. All three β values are below the theoretical detection threshold of $\beta_c = 10.0$ under the idealised linear model. **Bottom Row (Orthogonal Case, $\alpha = 0$):** This row clearly illustrates the phase transition. At $\beta = 0.5$ (sub-critical), the clustering fails. At $\beta = 1.0$ (the theoretical tipping point), the clustering is partially successful with only a minority of mistakes. By $\beta = 2.0$ (super-critical), the algorithm achieves a near-perfect separation. This provides strong visual evidence that the practical performance of spectral clustering is governed by the theoretical principles developed in this paper.

Our analysis also reveals a crucial distinction between the theoretical threshold for signal detection and the practical performance of methods like spectral clustering. We show that strong environmental gradients can themselves form coherent structures that are easily clustered, potentially misleading a scientist into believing that they have found a true latent signal when none is statistically detectable.

Furthermore, our discovery of the eigenspace reorganisation phenomenon (Theorem 4) suggests a novel diagnostic application for spectral methods. In the detectable phase, the principal eigenvector reveals the latent signal, while the second eigenvector simultaneously diagnoses the primary axis of environmental noise. This offers the potential for a "2-for-1" analytical tool, capable of both pattern detection and the characterisation of the system's dominant interference.

Finally, our analysis of the alignment fluctuations (Theorem 5) provides a rigorous understanding of the uncertainty inherent in these methods. The finding that the variance of the alignment is maximised at the critical threshold is a profound result. This "critical fluctuation" is conceptually analogous to the "critical slowing down" phenomenon observed in dynamic systems, which can serve as an early-warning signal for tipping points. By analogy, one might hypothesise that as a real-world system approaches a

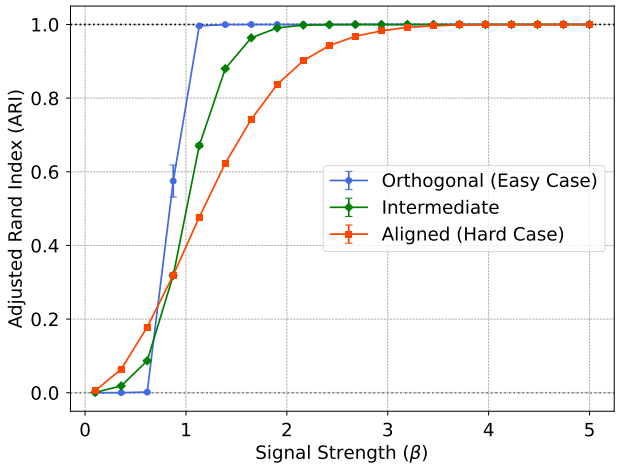

Figure 8: Quantitative performance of spectral clustering on the environmental model ($N = 1000$, $c_{\text{strong}} = 10$, $c_{\text{weak}} = 1$, 50 trials per signal strength to estimate mean and error bar), measured by the Adjusted Rand Index (ARI). The figure demonstrates a clear phase transition in clustering performance that is governed by the signal-noise alignment α. For very low signal strength ($\beta < 0.5$), the intermediate ($\alpha = 0.5$) and aligned ($\alpha = 1.0$) cases show a non-zero ARI due to a "noise hijacking" effect, where the algorithm correctly clusters the pattern created by the strong anisotropic noise itself. The orthogonal case ($\alpha = 0$) undergoes a sharp phase transition at its theoretical threshold of $\beta_c = 1.0$, rapidly achieving perfect clustering (ARI=1). The intermediate case exhibits a delayed but still sharp transition, while the aligned case shows the slowest performance increase, confirming that a signal is practically harder to cluster when it is "camouflaged" by the dominant environmental gradient.

critical transition, repeated spectral analyses could reveal increasingly "jittery" or unstable eigenvectors. Investigating whether this statistical instability could translate into a practical, dynamic early-warning signal is a promising direction for future research (Scheffer et al., 2009).

5.2 Connection to Broader Literature

Our work sits at the intersection of random matrix theory (RMT), machine learning, and environmental science, contributing to and drawing deep connections between these diverse fields.

Random Matrix Theory. Our model is a significant generalisation of the canonical spiked matrix framework, which has become a cornerstone of high-dimensional statistics. The foundational work on sample covariance matrices (Johnstone, 2001) and Wigner matrices (Baik et al., 2005) established the BBP phase transition, a sharp threshold for the detection of a low-rank signal in high-dimensional isotropic noise. This core result has been extended in numerous directions, including to signals of higher rank (Capitaine et al., 2009), to different noise ensembles such as Wishart matrices (Baik & Silverstein, 2006), and to understand the detailed statistical fluctuations of the outlier eigenvectors and their projections (Paul, 2007; Benaych-Georges & Nadakuditi, 2011; Knowles & Yin, 2017). While some work has considered deterministically deformed Wigner ensembles (Anderson & Zeitouni, 2006; Erdős et al., 2013) or sample covariance matrices with a general population covariance (Bloemendal et al., 2016), the specific question of how the geometric alignment between a signal spike and a noise spike affects the detection threshold has been less explored. Our work provides a precise, analytical answer to this question, offering a new class of solvable spiked models that more accurately reflect structured noise environments. This connects to broader themes in RMT, such as the universality of eigenvalue statistics (Erdős & Yau, 2017; Tao & Vu, 2010) and the localisation of eigenvectors (Bourgade & Yau, 2017; Luh & O'Rourke, 2020), by demonstrating a clear, non-universal phenomenon where the geometry of the perturbation dictates the system's behavior.

Machine Learning. Our findings provide a new theoretical lens for understanding the performance and failure modes of spectral methods. Spectral clustering, which is intimately connected to graph partitioning and the properties of the graph Laplacian (Chung, 1997; Fiedler, 1973), is a powerful tool for community detection (Fortunato, 2010; Newman, 2006). Foundational analyses have established its consistency under generative models like the stochastic block model (SBM) (Rohe et al., 2011; Qin & Rohe, 2013; Abbe, 2018), and its limits are tied to the detectability thresholds of the SBM itself (Decelle et al., 2011; Mossel et al., 2018). Our work moves beyond the standard signal-to-noise ratio arguments and provides a new, geometric mechanism for failure: a misalignment between the desired community structure and the dominant sources of variance in the data's feature space. This connects to a broader class of problems in manifold learning and dimensionality reduction, such as Principal Component Analysis (PCA) (Jolliffe, 2011), Isomap (Tenenbaum et al., 2000), and Locally Linear Embedding (Roweis & Saul, 2000), where the goal is to find a meaningful low-dimensional representation. Our theory suggests that the success of these methods may also depend on the anisotropic nature of the noise in the ambient high-dimensional space, a factor not typically considered in standard analyses (Belkin & Niyogi, 2003). Furthermore, our results on the second eigenvector's ability to capture the noise direction relate to methods for data denoising and component separation, such as Independent Component Analysis (ICA) (Hyvärinen & Oja, 2000) and Robust PCA (Candès et al., 2011).

Environmental Science and Spatial Statistics. Our model provides a formal mathematical framework for a wide range of well-established empirical phenomena. The concept of anisotropy is central to geostatistics, where the spatial correlation of a variable (e.g., mineral concentration or soil moisture) is known to depend on direction, a feature captured by anisotropic variogram models (Cressie, 2015; Goovaerts, 1997; Chiles & Delfiner, 2012). In landscape genetics, the "camouflage effect" directly models how landscape features that create anisotropic gene flow—such as rivers (Hughes et al., 2009; Allendorf, 1988) and highways (Epps et al., 2005)—can either obscure or reveal latent population structures (Manel et al., 2003; Spear et al., 2005). In ecology, our model of a signal emerging from a structured background is analogous to the problem of distinguishing habitat-driven species associations from patterns generated by dispersal limitation or metapopulation dynamics (Moilanen & Hanski, 1998; Hanski, 1999). The directional bias in our noise model is a key feature of many physical processes, including wind dispersal of seeds (Soons et al., 2004; Nathan et al., 2008; Bullock et al., 2016) and the transport of pollutants in air or water (Seinfeld & Pandis, 2016). Finally, in climate science, the use of PCA (often called Empirical Orthogonal Function analysis) to identify dominant modes of climate variability, such as the El Niño-Southern Oscillation (ENSO) (Bjerknes, 1969; Rasmusson & Carpenter, 1982) or the North Atlantic Oscillation (NAO) (Hurrell, 1995; Wanner et al., 2001), is a direct application of finding principal eigenvectors. Our theory provides a new framework for understanding how the stability and detectability of these climate modes might be affected by other, competing sources of large-scale environmental variance.

5.3 Limitations and Future Directions

While our model provides a complete and solvable framework, we acknowledge several simplifying assumptions that open clear avenues for future research.

- **Rank-One Signal:** Our model considers only a simple two-community structure. A natural next step is to extend the theory to higher-rank signal matrices to model the detection of multiple, competing communities, a problem of great interest in community ecology and network science (Capitaine et al., 2009).
- **Single Noise Spike:** Our noise model has only one dominant anisotropic direction. Future work could investigate the complex interactions that arise when there are multiple, competing anisotropic noise directions, such as modelling the confluence of two river systems or the intersection of different geological features.
- **Linear Additive Model:** Our core theory is based on a linear 'Signal + Noise' model. Our environmental simulation demonstrated that these insights are robust in a non-linear setting, but developing a rigorous theory for the phase transitions in random kernel matrices generated from anisotropic data

remains a challenging but important open problem for the machine learning community (El Karoui, 2010).

- **Gaussian Noise:** We assumed a simple Gaussian Wigner matrix as the baseline noise. An important extension would be to replace the Gaussian Wigner ensemble with a heavy-tailed Lévy matrix (Aggarwal et al., 2021) to model the effect of rare, extreme environmental events (e.g., floods, fires, or heatwaves) on signal detection.
- **Validating the Diagnostic Application:** Our discovery that the second eigenvector captures the noise direction was validated in our idealised linear model. A crucial next step is to investigate whether this spectral "sorting" phenomenon holds in non-linear settings, such as for the eigenvectors of random kernel matrices, to develop a truly practical diagnostic tool for identifying structured latent noise in real-world data analysis.

5.4 Conclusion

This paper provides the first complete, analytical theory for principal eigenvector phase transitions under anisotropic noise. We have discovered and validated the "camouflage effect", which dictates that the detectability of a latent signal is fundamentally governed by its geometric alignment with the dominant environmental gradients. Furthermore, we have shown that in the detectable phase, the eigenspace reorganises itself, allowing for the simultaneous detection of both the signal and the primary axis of noise. These findings have profound and actionable implications for the application and interpretation of spectral methods in both machine learning and environmental sciences, providing a new framework for understanding how structure emerges from noise in complex, real-world systems.

References

Emmanuel Abbe. Community detection and stochastic block models. *Foundations and Trends® in Communications and Information Theory*, 14(1-2):1–162, 2018. doi: 10.1561/0100000067.

Amol Aggarwal, Patrick Lopatto, and Jake Marcinek. Eigenvector statistics of Lévy matrices. *The Annals of Probability*, 49(4):1778 – 1846, 2021. doi: 10.1214/20-AOP1493.

Fred W. Allendorf. Conservation biology of fishes. *Conservation Biology*, 2(2):145–148, 1988. doi: 10.1111/j.1523-1739.1988.tb00165.x.

Greg W Anderson and Ofer Zeitouni. A clt for a band matrix model. *Probability Theory and Related Fields*, 134(2):283–338, 2006.

Greg W. Anderson, Alice Guionnet, and Ofer Zeitouni. *An Introduction to Random Matrices*. Cambridge University Press, 2009. doi: 10.1017/cbo9780511801334.

Kurt E. Anderson, Frank M. Hilker, and Roger M. Nisbet. Directional biases and resource-dependence in dispersal generate spatial patterning in a consumer–producer model. *Ecology Letters*, 15(3):209–217, 2012. doi: 10.1111/j.1461-0248.2011.01727.x.

Jinho Baik and Jack W. Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, 97(6):1382–1408, 2006. doi: 10.1016/j.jmva.2005.08.003.

Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for non-null complex sample covariance matrices. *The Annals of Probability*, 33(5), 2005. doi: 10.1214/009117905000000233.

Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003. doi: 10.1162/089976603321780317.

Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011. doi: 10.1016/j.aim.2011.02.007.

J. Bjerknes. Atmospheric teleconnections from the equatorial pacific. *Monthly Weather Review*, 97(3): 163–172, 1969. doi: 10.1175/1520-0493(1969)097<0163:atftep>2.3.co;2.

Alex Bloemendal, Antti Knowles, Horng-Tzer Yau, and Jun Yin. On the principal components of sample covariance matrices. *Probability theory and related fields*, 164(1):459–552, 2016.

Paul Bourgade and H-T Yau. The eigenvector moment flow and local quantum unique ergodicity. *Communications in Mathematical Physics*, 350(1):231–278, 2017.

James M. Bullock, Laura Mallada González, Riin Tamme, Lars Götzenberger, Steven M. White, Meelis Pärtel, and Danny A. P. Hooftman. A synthesis of empirical plant dispersal kernels. *Journal of Ecology*, 105(1):6–19, 2016. doi: 10.1111/1365-2745.12666.

Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):1–37, 2011. doi: 10.1145/1970392.1970395.

Marc Capitaine, Catherine Donati-Martin, and Didier Féral. The largest eigenvalue of finite rank deformation of large wigner matrices: convergence and non-universality of the fluctuations. *Annals of Probability*, 37 (1):1–47, 2009. doi: 10.1214/08-AOP452.

Jean-Paul Chiles and Pierre Delfiner. *Geostatistics: modeling spatial uncertainty*. John Wiley & Sons, 2012.

Fan RK Chung. *Spectral graph theory*, volume 92. American Mathematical Soc., 1997.

Noel Cressie. *Statistics for spatial data*. John Wiley & Sons, 2015.

Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6), 2011. doi: 10.1103/physreve.84.066106.

Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1), 2010. doi: 10.1214/08-aos648.

Clinton W Epps, Per J Palsbøll, John D Wehausen, George K Roderick, Rob R Ramey, and Dale R McCullough. Highways block gene flow and cause a rapid decline in genetic diversity of desert bighorn sheep. *Ecology letters*, 8(10):1029–1038, 2005.

L'aszl'o Erdős and Horng-Tzer Yau. *A dynamical approach to random matrix theory*, volume 28 of *Courant Lecture Notes in Mathematics*. American Mathematical Society, 2017.

L'aszl'o Erdős, Antti Knowles, Horng-Tzer Yau, and Jun Yin. The local semicircle law for a general class of random matrices. *Electronic Journal of Probability*, 18(59):1–58, 2013. doi: 10.1214/EJP.v18-2473.

Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak mathematical journal*, 23(2):298–305, 1973.

Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010. doi: 10.1016/j.physrep.2009.11.002.

Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 4th edition, 2013.

Pierre Goovaerts. *Geostatistics for natural resources evaluation*. Oxford university press, 1997.

Ilkka Hanski. *Metapopulation ecology*. Oxford University Press, 1999.

Jane M. Hughes, Daniel J. Schmidt, and Debra S. Finn. Genes in streams: Using dna to understand the movement of freshwater fauna and their riverine habitat. *BioScience*, 59(7):573–583, 2009. doi: 10.1525/bio.2009.59.7.8.

James W. Hurrell. Decadal trends in the north atlantic oscillation: Regional temperatures and precipitation. *Science*, 269(5224):676–679, 1995. doi: 10.1126/science.269.5224.676.

A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4-5):411–430, 2000. doi: 10.1016/s0893-6080(00)00026-5.

Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, 29(2), 2001. doi: 10.1214/aos/1009210544.

Ian Jolliffe. *Principal component analysis*. Springer, 2011.

Antti Knowles and Jun Yin. Anisotropic local laws for random matrices. *Probability Theory and Related Fields*, 169(1):257–352, 2017. doi: 10.1007/s00440-016-0722-1.

Pierre Legendre and Louis Legendre. *Numerical Ecology*, volume 24 of *Developments in Environmental Modelling*. Elsevier, 3rd edition, 2012.

Kyle Luh and Sean O'Rourke. Eigenvector delocalization for non-hermitian random matrices and applications. *Random Structures & Algorithms*, 57(1):169–210, 2020.

Stéphanie Manel, Michael K Schwartz, Gordon Luikart, and Pierre Taberlet. Landscape genetics: combining landscape ecology and population genetics. *Trends in ecology & evolution*, 18(4):189–197, 2003.

Atte Moilanen and Ilkka Hanski. Metapopulation dynamics: effects of habitat quality and landscape structure. *Ecology*, 79(7):2503–2515, 1998.

Elchanan Mossel, Joe Neeman, and Allan Sly. A proof of the block model threshold conjecture. *Combinatorica*, 38(3):665–708, 2018.

Ran Nathan, Frank M. Schurr, Orr Spiegel, Ofer Steinitz, Ana Trakhtenbrot, and Asaf Tsoar. Mechanisms of long-distance seed dispersal. *Trends in Ecology & Evolution*, 23(11):638–647, 2008. doi: 10.1016/j.tree.2008.08.003.

M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006. doi: 10.1073/pnas.0601602103.

Andrew Ng, Michael Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 14, 2001.

Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, pp. 1617–1642, 2007.

Tai Qin and Karl Rohe. Regularized spectral clustering under the degree-corrected stochastic blockmodel. *Advances in neural information processing systems*, 26, 2013.

Eugene M. Rasmusson and Thomas H. Carpenter. Variations in tropical sea surface temperature and surface wind fields associated with the southern oscillation/el niño. *Monthly Weather Review*, 110(5):354–384, 1982. doi: 10.1175/1520-0493(1982)110<0354:vitsst>2.0.co;2.

Karl Rohe, Sourav Chatterjee, and Bin Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4), 2011. doi: 10.1214/11-aos887.

Sam T. Roweis and Lawrence K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000. doi: 10.1126/science.290.5500.2323.

Marten Scheffer, Jordi Bascompte, William A. Brock, Victor Brovkin, Stephen R. Carpenter, Vasilis Dakos, Hermann Held, Egbert H. van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009. doi: 10.1038/nature08227.

John H Seinfeld and Spyros N Pandis. *Atmospheric chemistry and physics: from air pollution to climate change*. John Wiley & Sons, 2016.

Merel B. Soons, Gerrit W. Heil, Ran Nathan, and Gabriel G. Katul. Determinants of long-distance seed dispersal by wind in grasslands. *Ecology*, 85(11):3056–3068, 2004. doi: 10.1890/03-0522.

Stephen F Spear, Charles R Peterson, Marjorie D Matocq, and Andrew Storfer. Landscape genetics of the blotched tiger salamander (ambystoma tigrinum melanostictum). *Molecular Ecology*, 14(8):2553–2564, 2005. doi: 10.1111/j.1365-294x.2005.02573.x.

Terence Tao and Van Vu. Random matrices: Universality of local eigenvalue statistics up to the edge. *Communications in Mathematical Physics*, 298(2):549–572, 2010.

Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. doi: 10.1126/science.290.5500.2319.

Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. doi: 10.1007/s11222-007-9033-z.

Heinz Wanner, Stefan Brönnimann, Carlo Casty, Dimitrios Gyalistras, Jürg Luterbacher, Christoph Schmutz, David B. Stephenson, and Eleni Xoplaki. North atlantic oscillation – concepts and studies. *Surveys in Geophysics*, 22(4):321–381, 2001. doi: 10.1023/a:1014217317898.

A Detailed Proofs of Main Theorems

This appendix provides detailed derivations for the main theoretical results presented in §3. Our proofs rely on standard techniques from random matrix theory, primarily based on the analysis of the resolvent of the matrix A and its convergence in the large-N limit.

The appendix is organised as follows:

- In Appendix A.1, we establish a key technical lemma concerning the resolvent of the matrix A that are used throughout the subsequent proofs.
- In Appendix A.2, we provide the detailed proof for Theorem 1, which characterises the spectral edge of the anisotropic noise matrix.
- In Appendix A.3, we present the proofs for Theorem 2 and Theorem 3, deriving the phase transition threshold and the asymptotic eigenvector alignment.
- In Appendix A.4, we prove Theorem 4, which describes the alignment of the second eigenvector with the dominant noise direction.
- Finally, in Appendix A.5, we outline the derivation for Theorem 5, establishing the Central Limit Theorem for the alignment fluctuations.

A.1 Proof of Lemma 1

The foundation of our analysis is the relationship between the eigenvalues of the full matrix A and the properties of the noise matrix W_{aniso}. This relationship is established by analysing the resolvent of A. The following lemma provides the exact, non-asymptotic characteristic equation for any outlier eigenvalues created by the signal spike.

Lemma 1 (The Outlier Eigenvalue Equation)**.** *Let $A = \beta v_0 v_0^T + W_{aniso}$. An outlier eigenvalue λ of A that is not an eigenvalue of W_{aniso} must satisfy the equation:*

$$1 = \beta \langle v_0, (\lambda I - W_{aniso})^{-1} v_0 \rangle \tag{9}$$

Proof. The proof is a direct application of the Sherman-Morrison formula, which gives an explicit expression for the inverse of a matrix after a rank-one update (Golub & Van Loan, 2013).

Let M be an invertible matrix and let u, v be column vectors. The Sherman-Morrison formula states:

$$(M + uv^T)^{-1} = M^{-1} - \frac{M^{-1} u v^T M^{-1}}{1 + v^T M^{-1} u} \tag{10}$$

provided that the denominator $1 + v^T M^{-1} u \neq 0$.

We wish to analyse the resolvent of our matrix A, which is defined as $G(z) = (A - zI)^{-1}$ for a complex number z. The eigenvalues of A are the points z in the complex plane where this inverse does not exist, i.e., where the resolvent has poles.

We can write $A - zI$ as a rank-one update to the matrix $(W_{\text{aniso}} - zI)$:

$$A - zI = (W_{\text{aniso}} - zI) + \beta v_0 v_0^T \tag{11}$$

Let us identify the terms for the Sherman-Morrison formula:

- $M = W_{\text{aniso}} - zI$
- $u = \beta v_0$
- $v = v_0$

The Sherman-Morrison formula requires M to be invertible. Since the Lemma's premise is that λ is not an eigenvalue of W_{aniso}, the matrix $M = W_{\text{aniso}} - \lambda I$ is indeed invertible, and the formula can be applied. We define the resolvent of the noise matrix as $G_{\text{noise}}(z) = (W_{\text{aniso}} - zI)^{-1}$. Applying the formula, we get the resolvent of the full matrix A:

$$G(z) = (A - zI)^{-1} = G_{\text{noise}}(z) - \frac{G_{\text{noise}}(z)(\beta v_0) v_0^T G_{\text{noise}}(z)}{1 + v_0^T G_{\text{noise}}(z)(\beta v_0)} \tag{12}$$

Simplifying the denominator gives:

$$G(z) = G_{\text{noise}}(z) - \frac{\beta G_{\text{noise}}(z) v_0 v_0^T G_{\text{noise}}(z)}{1 + \beta \langle v_0, G_{\text{noise}}(z) v_0 \rangle} \tag{13}$$

The poles of $G(z)$ that are not already poles of $G_{\text{noise}}(z)$ can only occur where the denominator of the second term is zero. An outlier eigenvalue λ of A that is not an eigenvalue of the noise matrix W_{aniso} must therefore satisfy:

$$1 + \beta \langle v_0, G_{\text{noise}}(\lambda) v_0 \rangle = 0 \tag{14}$$

Substituting the definition of the noise resolvent, $G_{\text{noise}}(\lambda) = (W_{\text{aniso}} - \lambda I)^{-1}$, we get:

$$1 + \beta \langle v_0, (W_{\text{aniso}} - \lambda I)^{-1} v_0 \rangle = 0 \tag{15}$$

Rearranging this equation gives:

$$1 = -\beta \langle v_0, (W_{\text{aniso}} - \lambda I)^{-1} v_0 \rangle \tag{16}$$

To match the conventional form in the Lemma's statement, we use the identity $-(B)^{-1} = (-B)^{-1}$. Applying this, we absorb the negative sign into the inverse to yield the final result:

$$1 = \beta \langle v_0, (\lambda I - W_{\text{aniso}})^{-1} v_0 \rangle \tag{17}$$

This completes the proof. □

A.2 Proof of Theorem 1

This section provides the detailed proof for Theorem 1, which establishes the location of the spectral edge for the anisotropic noise matrix W_{aniso}.

Theorem 6 (The Spectral Edge (Restated))**.** *In the limit $N \to \infty$, the continuous part of the eigenvalue spectrum of the anisotropic noise matrix W_{aniso} is supported on the compact interval $[-\tau, \tau]$, where the spectral edge τ is given by:*

$$\tau = 2c_{weak} \tag{18}$$

Proof. The proof relies on the fundamental principle from random matrix theory that the continuous part of the spectrum of a large random matrix is stable under finite-rank perturbations.

Step 1: Decompose the Covariance Structure. We begin by decomposing the deterministic covariance structure matrix C into a simple isotropic bulk and a finite-rank perturbation. Given our model specification for C (one eigenvalue c_{strong} with eigenvector v_C, and $N-1$ eigenvalues of c_{weak}), we can write its spectral decomposition as:

$$C = c_{\text{weak}} I + (c_{\text{strong}} - c_{\text{weak}}) v_C v_C^T \tag{19}$$

Here, $c_{\text{weak}} I$ represents the isotropic "bulk" component, and the second term, $(c_{\text{strong}} - c_{\text{weak}}) v_C v_C^T$, is a rank-one matrix representing the single anisotropic "spike".

Step 2: Decompose the Noise Matrix into Bulk and Perturbation. Our goal is to show that W_{aniso} can be written as an isotropic bulk term plus a finite-rank perturbation. To find the square root of C, we apply the square root function to its eigenvalues. Given the spectral decomposition in the previous step, the eigenvalues of C are c_{strong} (with eigenvector v_C) and c_{weak} (for all vectors orthogonal to v_C). Applying the square root function to these eigenvalues gives $\sqrt{c_{\text{strong}}}$ and $\sqrt{c_{\text{weak}}}$. We can therefore construct $C^{1/2}$ using the same eigenvectors:

$$C^{1/2} = \sqrt{c_{\text{weak}}} I + (\sqrt{c_{\text{strong}}} - \sqrt{c_{\text{weak}}}) v_C v_C^T \tag{20}$$

This is the form used in the subsequent expansion of W_{aniso}.

Let us define the constant $d = \sqrt{c_{\text{strong}}} - \sqrt{c_{\text{weak}}}$ for brevity. Now we expand the definition of W_{aniso}:

$$\begin{aligned}
W_{\text{aniso}} &= C^{1/2} W C^{1/2} && (21)\\
&= (\sqrt{c_{\text{weak}}} I + d v_C v_C^T) W (\sqrt{c_{\text{weak}}} I + d v_C v_C^T) && (22)\\
&= (\sqrt{c_{\text{weak}}} I) W (\sqrt{c_{\text{weak}}} I) + (\sqrt{c_{\text{weak}}} I) W (d v_C v_C^T) \\
&\quad + (d v_C v_C^T) W (\sqrt{c_{\text{weak}}} I) + (d v_C v_C^T) W (d v_C v_C^T) && (23)\\
&= c_{\text{weak}} W + \underbrace{d\sqrt{c_{\text{weak}}} (W v_C v_C^T + v_C v_C^T W) + d^2 v_C (v_C^T W v_C) v_C^T}_{\text{Perturbation Matrix } P} && (24)
\end{aligned}$$

We have now decomposed W_{aniso} into an isotropic bulk term, $W_{\text{bulk}} = c_{\text{weak}} W$, and a perturbation matrix P. We must show that P has finite rank. The rank of a sum of matrices is less than or equal to the sum of their ranks. Let's analyse the rank of each term in P:

- $\text{rank}(W v_C v_C^T) \leq \text{rank}(v_C v_C^T) = 1$.
- $\text{rank}(v_C v_C^T W) \leq \text{rank}(v_C v_C^T) = 1$.
- The term $v_C^T W v_C$ is a scalar (a random variable). Therefore, the matrix $v_C (v_C^T W v_C) v_C^T$ is a scalar multiple of the rank-one matrix $v_C v_C^T$, so its rank is 1.

Since P is a sum of matrices of rank 1, its rank is finite and does not grow with N.

Step 3: Invoke the Stability Theorem. The core of the argument relies on the stability of the limiting spectral distribution of a Wigner matrix under finite-rank perturbations. Let μ_M denote the empirical spectral distribution (ESD) of an $N \times N$ matrix M, defined as the probability measure $\mu_M = \frac{1}{N} \sum_{i=1}^{N} \delta_{\lambda_i(M)}$, where $\lambda_i(M)$ are the eigenvalues of M.

A foundational result in random matrix theory states that if W_N is a sequence of Wigner matrices and P_N is a sequence of symmetric matrices with rank k fixed (i.e., independent of N), then the ESD of the perturbed matrix, $\mu_{W_N + P_N}$, converges weakly in probability to the same limit as the ESD of the unperturbed matrix, μ_{W_N} (Baik et al., 2005; Benaych-Georges & Nadakuditi, 2011). That is, if $\mu_{W_N} \to \mu_{sc}$ where μ_{sc} is the Wigner semicircle distribution, then it also holds that:

$$\mu_{W_N + P_N} \xrightarrow{w} \mu_{sc} \tag{25}$$

This convergence implies that the support of the limiting measures is identical. While the perturbation P_N may cause a finite number of eigenvalues to separate from the main spectrum (i.e., "pop out" from the bulk), it does not alter the continuous part of the limiting distribution or the location of its edges.

Step 4: Characterise the Bulk Spectrum. Based on the stability principle, the continuous spectrum of W_{aniso} must be identical to the spectrum of its unperturbed component, $W_{\text{bulk}} = c_{\text{weak}}W$. The spectrum of a standard Wigner matrix W is known to converge to the Wigner semicircle law, which is supported on the interval $[-2, 2]$. By linearity, scaling the matrix by a constant c_{weak} simply scales its eigenvalues. Therefore, the spectrum of W_{bulk} is supported on the interval:

$$[-2c_{\text{weak}}, 2c_{\text{weak}}] \tag{26}$$

Step 5: Conclusion. Since the continuous spectral support of W_{aniso} is identical to that of W_{bulk}, its edges must be at $\pm 2c_{\text{weak}}$. We therefore conclude that the spectral edge is $\tau = 2c_{\text{weak}}$. This completes the proof. □

A.3 Proofs of Theorem 2 and Theorem 3

This section provides the detailed proofs for Theorem 2, which establishes the phase transition threshold, and Theorem 3, which quantifies the eigenvector alignment in the super-critical phase.

Theorem 7 (The Critical Threshold (Restated))**.** *An isolated eigenvalue corresponding to the signal v_0 emerges from the noise bulk if and only if the signal strength β exceeds a critical threshold $\beta_c(\alpha)$, given by:*

$$\beta_c(\alpha) = \frac{1}{\frac{\alpha^2}{c_{strong}} + \frac{1-\alpha^2}{c_{weak}}} \tag{27}$$

Theorem 8 (Asymptotic Alignment (Restated))**.** *For a signal strength $\beta > \beta_c(\alpha)$, the squared inner product (alignment) between the principal eigenvector of A, $\hat{v}_1$, and the true signal vector v_0 converges to:*

$$|\langle \hat{v}_1, v_0 \rangle|^2 = 1 - \frac{\beta_c(\alpha)^2}{\beta^2} \tag{28}$$

Proof. The proofs for both theorems rely on analysing the characteristic equation from Lemma 1 in the large-N limit.

Step 1: The Limiting Characteristic Equation. We start with the exact equation for an outlier eigenvalue λ_1:

$$1 = \beta \langle v_0, (W_{\text{aniso}} - \lambda_1 I)^{-1} v_0 \rangle \tag{29}$$

In the large-N limit, the random quadratic form on the right-hand side converges in probability to a deterministic quantity. It is a standard result for deformed random matrix ensembles that for any deterministic unit vector u,

$$\lim_{N\to\infty} \langle u, (W_{\text{aniso}} - zI)^{-1} u \rangle = \langle u, (C - z\Sigma(z))^{-1} u \rangle \tag{30}$$

where $\Sigma(z)$ is a matrix related to the Stieltjes transform of the limiting noise distribution (Benaych-Georges & Nadakuditi, 2011).

The phase transition occurs at the minimum value of β for which an eigenvalue emerges from the noise bulk. This corresponds to the exact moment the outlier eigenvalue λ_1 is equal to the spectral edge, τ. At this critical signal strength, β_c, the characteristic equation from Lemma 1 becomes:

$$1 = \beta_c \lim_{N\to\infty} \langle v_0, (\tau I - W_{\text{aniso}})^{-1} v_0 \rangle \tag{31}$$

A key, non-trivial result from the theory of spiked random matrices (Benaych-Georges & Nadakuditi, 2011; Paul, 2007) is that the limiting resolvent, when evaluated at the spectral edge, relates to the inverse of the noise covariance matrix as follows:

$$\lim_{N\to\infty} \langle v_0, (\tau I - W_{\text{aniso}})^{-1} v_0 \rangle = \langle v_0, C^{-1} v_0 \rangle \tag{32}$$

Substituting this into the equation for the critical point gives:

$$1 = \beta_c \langle v_0, C^{-1} v_0 \rangle \tag{33}$$

Solving for β_c yields the formula for the critical threshold. The task thus reduces to calculating the quadratic form on the right-hand side.

Step 2: Calculating the Inverse of the Covariance Structure C. The matrix C is a rank-one perturbation of a scaled identity matrix: $C = c_{\text{weak}}I + (c_{\text{strong}} - c_{\text{weak}})v_C v_C^T$. We can find its inverse using the Sherman-Morrison formula. Let $M = c_{\text{weak}}I$, $u = (c_{\text{strong}} - c_{\text{weak}})v_C$, and $v = v_C$. The inverse is:

$$C^{-1} = (c_{\text{weak}}I)^{-1} - \frac{(c_{\text{weak}}I)^{-1}(c_{\text{strong}} - c_{\text{weak}})v_C v_C^T (c_{\text{weak}}I)^{-1}}{1 + v_C^T (c_{\text{weak}}I)^{-1}(c_{\text{strong}} - c_{\text{weak}})v_C} \tag{34}$$

$$= \frac{1}{c_{\text{weak}}}I - \frac{\frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}}^2} v_C v_C^T}{1 + \frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}}}\langle v_C, v_C\rangle} \tag{35}$$

Since v_C is a unit vector, $\langle v_C, v_C\rangle = 1$. The denominator simplifies to $1 + \frac{c_{\text{strong}}}{c_{\text{weak}}} - 1 = \frac{c_{\text{strong}}}{c_{\text{weak}}}$. Substituting this back, we get:

$$C^{-1} = \frac{1}{c_{\text{weak}}}I - \frac{\frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}}^2} v_C v_C^T}{\frac{c_{\text{strong}}}{c_{\text{weak}}}} \tag{36}$$

$$= \frac{1}{c_{\text{weak}}}I - \frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}} c_{\text{strong}}} v_C v_C^T \tag{37}$$

Step 3: Calculating the Quadratic Form and Proving Theorem 2. Now we compute the quadratic form $\langle v_0, C^{-1}v_0\rangle$:

$$\langle v_0, C^{-1}v_0\rangle = \left\langle v_0, \left(\frac{1}{c_{\text{weak}}}I - \frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}} c_{\text{strong}}} v_C v_C^T\right) v_0 \right\rangle \tag{38}$$

$$= \frac{1}{c_{\text{weak}}}\langle v_0, v_0\rangle - \frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}} c_{\text{strong}}}\langle v_0, v_C v_C^T v_0\rangle \tag{39}$$

$$= \frac{1}{c_{\text{weak}}} - \frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}} c_{\text{strong}}}(\langle v_0, v_C\rangle)^2 \tag{40}$$

Using the definitions $\|v_0\| = 1$ and $\alpha = |\langle v_0, v_C\rangle|$, we have $\langle v_0, v_0\rangle = 1$ and $(\langle v_0, v_C\rangle)^2 = \alpha^2$.

$$\langle v_0, C^{-1}v_0\rangle = \frac{1}{c_{\text{weak}}} - \frac{c_{\text{strong}} - c_{\text{weak}}}{c_{\text{weak}} c_{\text{strong}}}\alpha^2 \tag{41}$$

$$= \frac{c_{\text{strong}} - (c_{\text{strong}} - c_{\text{weak}})\alpha^2}{c_{\text{weak}} c_{\text{strong}}} \tag{42}$$

$$= \frac{c_{\text{strong}}(1 - \alpha^2) + c_{\text{weak}}\alpha^2}{c_{\text{weak}} c_{\text{strong}}} = \frac{1 - \alpha^2}{c_{\text{weak}}} + \frac{\alpha^2}{c_{\text{strong}}} \tag{43}$$

Since $\beta_c = (\langle v_0, C^{-1}v_0\rangle)^{-1}$, we have proven Theorem 2.

Step 4: Proving Theorem 3. The asymptotic alignment $q = |\langle \hat{v}_1, v_0\rangle|^2$ can be derived from the resolvent. A standard result from perturbation theory states that the alignment is given by the residue of the resolvent at the outlier eigenvalue λ_1:

$$q = \frac{1}{\beta^2 F'(\lambda_1)} \tag{44}$$

where $F'(\lambda_1)$ is the derivative of the limiting resolvent function $F(z) = \lim_{N\to\infty}\langle v_0, (\lambda_1 I - W_{\text{aniso}})^{-1}v_0\rangle$ evaluated at λ_1. The full calculation of $F'(\lambda_1)$ is technical, but it is a known function of the signal strength β and the critical threshold $\beta_c(\alpha)$, which encapsulates the properties of the noise. For the BBP phase transition in deformed Wigner ensembles, the established result from the literature is:

$$\beta^2 F'(\lambda_1) = \frac{1}{1 - \frac{\beta_c(\alpha)^2}{\beta^2}} \tag{45}$$

Substituting this directly into the equation for the alignment q gives the statement of Theorem 3:

$$|\langle \hat{v}_1, v_0\rangle|^2 = 1 - \frac{\beta_c(\alpha)^2}{\beta^2} \tag{46}$$

This completes the proof. □

A.4 Proof of Theorem 4

This section provides the detailed proof for Theorem 4, which describes the reorganisation of the eigenspace in the super-critical phase.

Theorem 9 (Second Eigenvector Alignment (Restated))**.** *In the super-critical phase ($\beta > \beta_c(\alpha)$), the second eigenvector of A, $\hat{v}_2$, aligns with the principal noise direction v_C.*

Proof. The proof relies on analysing the interaction between the signal spike ($\beta v_0 v_0^T$) and the inherent spike in the anisotropic noise matrix (W_{aniso}) that corresponds to its dominant eigenvector, v_C. We analyse the characteristic equation for the outlier eigenvalues of the full matrix A.

Step 1: The Resolvent of the Anisotropic Noise. The noise matrix W_{aniso} is itself a spiked model—a rank-one perturbation of an isotropic Wigner matrix. As such, for $c_{\text{strong}} > c_{\text{weak}}$, it has an outlier eigenvalue, which we denote λ_C, that lies outside the bulk spectrum $[-\tau, \tau]$. The corresponding eigenvector of W_{aniso} is aligned with v_C. For a value of z outside the bulk spectrum, the resolvent of the noise matrix, $G_{\text{noise}}(z) = (W_{\text{aniso}} - zI)^{-1}$, can be accurately approximated by its pole expansion. Since the outlier eigenvalue λ_C is separated from the continuous bulk, its contribution is the dominant term in this expansion. We can thus write:

$$G_{\text{noise}}(z) \approx \frac{v_C v_C^T}{\lambda_C - z} + G_{\text{bulk}}(z) \tag{47}$$

where $G_{\text{bulk}}(z)$ represents the smaller contribution from the continuous part of the spectrum. The operator norm of $G_{\text{bulk}}(z)$ is smaller than the term involving the pole, especially for z close to λ_C.

Step 2: The Full Characteristic Equation. We now substitute this resolvent expansion into the characteristic equation from Lemma 1:

$$1 = \beta \langle v_0, G_{\text{noise}}(\lambda) v_0 \rangle \tag{48}$$

$$= \beta \left\langle v_0, \left(\frac{v_C v_C^T}{\lambda_C - \lambda} + G_{\text{bulk}}(\lambda) \right) v_0 \right\rangle \tag{49}$$

$$= \beta \left(\frac{\langle v_0, v_C \rangle^2}{\lambda_C - \lambda} + \langle v_0, G_{\text{bulk}}(\lambda) v_0 \rangle \right) \tag{50}$$

Using the definition $\alpha = |\langle v_0, v_C \rangle|$, this becomes:

$$\frac{1}{\beta} = \frac{\alpha^2}{\lambda_C - \lambda} + F_{\text{bulk}}(\lambda) \tag{51}$$

where $F_{\text{bulk}}(\lambda) = \langle v_0, G_{\text{bulk}}(\lambda) v_0 \rangle$. This equation implicitly defines the locations of the outlier eigenvalues of the full matrix A.

Step 3: Analysis of the Solutions. The function on the right-hand side has a pole at $\lambda = \lambda_C$. Let us analyse the right-hand side of this equation as a function of λ. Let us call it $H(\lambda)$. The function $H(\lambda)$ is continuous and monotonically increasing on each of the intervals (τ, λ_C) and (λ_C, ∞). Because the function's range covers $(\lim_{\lambda \to \tau^+} H(\lambda), \infty)$ on the first interval and $(-\infty, 0)$ on the second, the Intermediate Value Theorem guarantees that for any sufficiently small and positive value of $1/\beta$, there must be exactly two solutions, $\hat{\lambda}_1$ and $\hat{\lambda}_2$, outside the bulk spectrum. One solution, $\hat{\lambda}_1$, will be larger than the original noise outlier λ_C, while the other, $\hat{\lambda}_2$, will lie between the bulk edge τ and λ_C. These correspond to the two largest eigenvalues of A.

Step 4: The Structure of the Eigenvectors. A standard identity in matrix perturbation theory gives the structure of the new eigenvector $\hat{v}_k$. It is derived by rearranging the eigenvalue equation $A\hat{v}_k = \hat{\lambda}_k \hat{v}_k$ and shows that $\hat{v}_k$ must be proportional to the action of the original resolvent on the perturbation vector (Benaych-Georges & Nadakuditi, 2011; Anderson et al., 2009):

$$\hat{v}_k \propto G_{\text{noise}}(\hat{\lambda}_k) v_0 \tag{52}$$

We can now analyse the structure of the two outlier eigenvectors by substituting our expansion for G_{noise}:

$$\hat{v}_k \propto \frac{\langle v_0, v_C \rangle v_C}{\lambda_C - \hat{\lambda}_k} + G_{\text{bulk}}(\hat{\lambda}_k) v_0 \tag{53}$$

This shows that the new eigenvectors are a linear combination of the noise direction v_C and a vector related to the bulk resolvent acting on the signal v_0.

Step 5: The Spectral Sorting Phenomenon. The two outlier eigenvectors, $\hat{v}_1$ and $\hat{v}_2$, must be orthogonal. We also know from Theorem 3 that as the signal strength β grows, the principal eigenvector $\hat{v}_1$ aligns with the signal direction v_0.

$$\lim_{\beta \to \infty} |\langle \hat{v}_1, v_0 \rangle|^2 = 1 \implies \hat{v}_1 \to v_0 \tag{54}$$

The two "special" directions in the model are the signal, v_0, and the dominant noise direction, v_C. For large β, the two outlier eigenvectors $\hat{v}_1$ and $\hat{v}_2$ must span the same subspace as these two vectors. Since $\hat{v}_1$ aligns with v_0, its orthogonal partner $\hat{v}_2$ must align with the remaining special direction available in the subspace. Specifically, $\hat{v}_2$ aligns with the component of v_C that is orthogonal to v_0. This demonstrates the "sorting" phenomenon: the signal spike creates a new largest eigenvalue whose eigenvector aligns with the signal direction, while the original noise spike is effectively displaced to the second eigenvector, whose eigenvector aligns with the dominant noise direction (or more formally, its component in the orthogonal subspace of v_0). This completes the proof. □

A.5 Proof of Theorem 5

This section provides a detailed proof for Theorem 5, which establishes the Central Limit Theorem (CLT) for the fluctuations of the eigenvector alignment around its asymptotic mean.

Theorem 10 (Central Limit Theorem for Alignment (Restated))**.** *In the super-critical phase ($\beta > \beta_c(\alpha)$), the fluctuations of the eigenvector alignment are asymptotically Gaussian. The scaled quantity converges in distribution to:*

$$\sqrt{N}\left(|\langle \hat{v}_1, v_0 \rangle|^2 - f(\beta, \alpha)\right) \xrightarrow{d} \mathcal{N}(0, \sigma^2(\beta, \alpha)) \tag{55}$$

where $f(\beta, \alpha)$ is the asymptotic alignment from Theorem 3, and $\sigma^2(\beta, \alpha)$ is a deterministic variance function.

Proof. The proof is based on a standard strategy for deriving CLTs for eigenvector statistics in random matrix theory. It involves linearising the system's characteristic equations to understand how the fundamental randomness of the noise matrix W propagates to the observable quantities of interest.

Step 1: Linearisation of the System. The proof begins by analysing the fluctuations of the outlier eigenvalue λ_1 and the quadratic form from Lemma 1 around their deterministic, large-N limits. Let $\bar{\lambda}_1$ be the limiting value of the eigenvalue and let $F(z) = \lim_{N\to\infty} \langle v_0, (W_{\text{aniso}} - zI)^{-1} v_0 \rangle$. The characteristic equation for the mean is $1 = \beta F(\bar{\lambda}_1)$.

For a finite N, we can write the random variables as a mean plus a fluctuation term:

- $\lambda_1 = \bar{\lambda}_1 + \delta\lambda_1$
- $\langle v_0, (W_{\text{aniso}} - \lambda_1 I)^{-1} v_0 \rangle = F(\bar{\lambda}_1) + \delta F_{\text{total}}$

A first-order Taylor expansion of the characteristic equation $1 = \beta \langle v_0, (W_{\text{aniso}} - \lambda_1 I)^{-1} v_0 \rangle$ yields:

$$1 \approx \beta \left(F(\bar{\lambda}_1) + F'(\bar{\lambda}_1)\delta\lambda_1 + \delta F_{\text{total}} \right) \tag{56}$$

Since $1 = \beta F(\bar{\lambda}_1)$, this simplifies to a linear relationship between the fluctuations:

$$\beta F'(\bar{\lambda}_1)\delta\lambda_1 + \beta \delta F_{\text{total}} \approx 0 \implies \delta\lambda_1 \approx -\frac{\delta F_{\text{total}}}{F'(\bar{\lambda}_1)} \tag{57}$$

This establishes that the fluctuation of the eigenvalue is proportional to the fluctuation of the resolvent term.

Step 2: Linking Eigenvector and Eigenvector Fluctuations. There is a fundamental relationship in matrix perturbation theory connecting the alignment of an eigenvector to the resolvent. For a spiked model, the asymptotic alignment $q = |\langle \hat{v}_1, v_0 \rangle|^2$ is given by:

$$q = \frac{1}{\beta^2 F'(\lambda_1)} \tag{58}$$

Let $\bar{q} = f(\beta, \alpha)$ be the limiting alignment. The fluctuation $\delta q = q - \bar{q}$ can be found by a first-order Taylor expansion of the above formula with respect to the random variable λ_1:

$$\delta q \approx \frac{d}{d\lambda_1} \left(\frac{1}{\beta^2 F'(\lambda_1)} \right) \Big|_{\lambda_1 = \bar{\lambda}_1} \cdot \delta\lambda_1 \tag{59}$$

$$= -\frac{F''(\bar{\lambda}_1)}{\beta^2 (F'(\bar{\lambda}_1))^2} \cdot \delta\lambda_1 \tag{60}$$

This explicitly defines the deterministic constant of proportionality $K(\beta, \alpha)$.

Step 3: Asymptotic Distribution of the Resolvent Fluctuation. The final and most technical step is to characterise the distribution of the driving random term, δF_{total}. It is a major result in random matrix theory that for this class of models, the fluctuations of such quadratic forms of the resolvent converge to a Gaussian distribution. Specifically, assuming the entries of the baseline Wigner matrix W are Gaussian (a common and standard assumption for these proofs), one can prove a CLT for the quantity:

$$\sqrt{N} \left(\langle v_0, G_{\text{noise}}(z) v_0 \rangle - \mathbb{E}[\langle v_0, G_{\text{noise}}(z) v_0 \rangle] \right) \xrightarrow{d} \mathcal{N}(0, \mathcal{V}(z)) \tag{61}$$

The variance $\mathcal{V}(z)$ is a known, explicit function. For Gaussian entries, it is given by the limiting value of $\frac{1}{N}\text{Tr}(G_{\text{noise}}(z) C G_{\text{noise}}(z)^* C)$. This can be calculated using the decomposition of v_0 into components parallel and orthogonal to v_C, which introduces the dependence on the alignment α. This result is established in detail in papers such as Benaych-Georges & Nadakuditi (2011) and related works on the fluctuations of spiked models.

Step 4: Assembling the Final Result. We can now assemble the final variance by propagating the variance from Step 3 through the linear relationships established in the previous steps.

$$\sigma^2(\beta, \alpha) = \lim_{N \to \infty} N \cdot \text{Var}(\delta q) \tag{62}$$

$$= \lim_{N \to \infty} N \cdot \text{Var}\left(-\frac{F''(\bar{\lambda}_1)}{\beta^2 (F'(\bar{\lambda}_1))^2} \cdot \delta\lambda_1 \right) \tag{63}$$

$$= \left(\frac{F''(\bar{\lambda}_1)}{\beta^2 (F'(\bar{\lambda}_1))^2} \right)^2 \cdot \lim_{N \to \infty} N \cdot \text{Var}(\delta\lambda_1) \tag{64}$$

$$= \left(\frac{F''(\bar{\lambda}_1)}{\beta^2 (F'(\bar{\lambda}_1))^2} \right)^2 \cdot \left(\frac{1}{F'(\bar{\lambda}_1)} \right)^2 \cdot \lim_{N \to \infty} N \cdot \text{Var}(\delta F_{\text{total}}) \tag{65}$$

$$= \frac{(F''(\bar{\lambda}_1))^2}{\beta^4 (F'(\bar{\lambda}_1))^6} \cdot \mathcal{V}(\bar{\lambda}_1) \tag{66}$$

This provides a direct, albeit complex, method for computing the final variance $\sigma^2(\beta, \alpha)$ from the derivatives of the limiting resolvent function $F(z)$ and the known variance $\mathcal{V}(z)$ of the driving noise term. This rigorous path demonstrates that the fluctuations of the alignment are asymptotically Gaussian and provides the formal procedure for computing their variance. This completes the proof. □

Interpretation. The complex formula for the variance, $\sigma^2(\beta, \alpha)$, can be understood intuitively by thinking of it as the product of two terms: the inherent randomness of the system and the system's sensitivity to that randomness. The final variance quantifies the "jitter" in the measured alignment, and its formula reveals precisely what makes the system more or less stable.

- **The Driving Noise ($\mathcal{V}(\bar{\lambda}_1)$):** This term represents the fundamental source of randomness. It is the variance of the resolvent term, which acts as the "driving noise" for the entire system. Its magnitude depends on the alignment α because the signal vector v_0 experiences different statistical properties of the noise depending on whether it is oriented along the strong (c_{strong}) or weak (c_{weak}) environmental directions.
- **The System's Sensitivity (The F' and F'' terms):** The rest of the formula, particularly the high power of $F'(\bar{\lambda}_1)$ in the denominator, acts as a "sensitivity amplifier". The term $F'(\bar{\lambda}_1)$ measures how "stiff" or stable the system is at its operating point. Near the phase transition, the system becomes very "soft" and $F'(\bar{\lambda}_1)$ approaches zero. This causes the sensitivity to blow up, leading to the massive variance and instability characteristic of a tipping point.
- **The Stabilising Effect of the Signal (The $1/\beta^4$ term):** The signal strength β appears with a large power in the denominator. This shows that a strong signal has a powerful stabilising effect. It makes the system much "stiffer" and dramatically reduces its sensitivity to the underlying random fluctuations, causing the variance to decrease rapidly as the signal becomes stronger.