Spectral Regime Detection: Change-Point Identification via Eigenmode Drift

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper

Abstract

We introduce spectral regime detection: a method to identify structural breaks in time series and panel data by monitoring the drift of spectral coefficients over sliding windows. The key insight is that a regime change is a **jump in the spectral representation of the data-generating process** — the coefficient vector $A = (A_1, \dots, A_K)$ moves discontinuously when the underlying regression relationship changes.

The method requires a **shared eigenbasis** across windows, computed once from a reference period. Each window's data is projected onto this basis, yielding comparable coefficient vectors. The drift $d_t = \|A_t - A_{t-1}\|$ between consecutive windows is monitored. Under a stable regime, drift is $O(\sigma/\sqrt{n_{window}})$; at a regime change, drift spikes to $O(\|\beta_{new} - \beta_{old}\|)$.

On synthetic data with known regime changes at $t = 500$ and $t = 800$ (structural break in regression coefficients), the detector identifies both changes with **20–40 sample lag** and **6–10x drift spike** relative to within-regime fluctuation. The method requires no distributional assumptions, works with any number of features, and naturally adapts to the signal-to-noise ratio through the spectral shrinkage filter.

---

1. Introduction

Regime detection — identifying when the data-generating process changes — is fundamental in finance (market regimes, policy shifts), economics (business cycles), and engineering (fault detection). Existing methods fall into two categories:

Parametric methods (Hamilton, 1989; Bai and Perron, 1998): assume a specific model (e.g., Markov switching, piecewise linear) and estimate change points within the model. Limitation: the regime model must be specified in advance.

Distribution-based methods (CUSUM, MOSUM, kernel change-point detection): monitor statistics of the data distribution. Limitation: high-dimensional data requires careful dimension reduction; the statistics are not directly interpretable.

We propose a third approach: **spectral regime detection**. Instead of monitoring distributional statistics or fitting a regime model, we monitor the **spectral coefficients** of a regression relationship over sliding windows. The coefficients are the projection of the response onto the data's eigenmodes — orthogonal directions of variation ordered by importance.

1.1 Why Spectral?

A regime change alters the relationship $y = f(X)+\epsilon$. In eigenspace, this means the coefficient vector $A = (A_1, \dots, A_K)$ changes. Because eigenmodes are orthogonal, each mode changes independently — you can see WHICH aspects of the relationship changed, not just that something changed.

Existing method	Detects	Explains
CUSUM	change in mean	no
Markov switching	regime probability	model-specific
Kernel change-point	distributional shift	no
Spectral regime	**coefficient jump**	**which modes changed**

1.2 Contributions

1. **Shared-basis spectral monitoring**: project each window onto a fixed eigenbasis for comparable coefficients
2. **Drift statistic**: $d_t = \|A_t - A_{t-1}\|$ with z-score normalization against rolling history
3. **Mode-level diagnosis**: not just "change detected" but "modes 1 and 3 jumped, mode 2 stable" — structural interpretation of what changed
4. **No distributional assumptions**: works for any data where regression is meaningful
5. **Experimental validation**: detection lag 20–40 samples, 6–10x drift spike, on synthetic benchmarks

---

2. Method

2.1 Shared Eigenbasis

The critical design choice: eigenvectors must be **shared** across windows. If each window computes its own eigendecomposition, the bases are incomparable and coefficients cannot be subtracted.

Construction: from a reference window X_{ref} (e.g., the first window or a burn-in period), compute the SVD:

$$X_{ref} = U\Sigma V^T$$

The top-K right singular vectors $V_K = [v_1, \dots, v_K]$ form the shared basis in feature space ($p \times K$).

For all subsequent windows t, the coefficients are:

$$A_t = (X_t V_K)^+ y_t$$

where $(X_t V_K)^+$ is the pseudoinverse. This projects the response onto the **same** K directions across all windows.

2.2 Drift Statistic

Between consecutive windows:

$$d_t = \|A_t - A_{t-1}\|_2$$

Under a stable regime with noise variance σ^2 and window size n_w:

$$\mathbb{E}[d_t] = O\left(\frac{\sigma}{\sqrt{n_w}} \cdot \sqrt{K}\right), \qquad \mathrm{Var}(d_t) = O\left(\frac{\sigma^2 K}{n_w}\right)$$

At a regime change where the true coefficient vector shifts from β_{old} to β_{new}:

$$d_t \approx \|V_K^T(\beta_{new} - \beta_{old})\|_2$$

The signal-to-noise ratio of the detector is:

$$\mathrm{SNR} = \frac{\|V_K^T \Delta\beta\|}{\sigma\sqrt{K/n_w}}$$

For $n_w = 80$, $K = 8$, $\sigma = 1.5$, and $\|\Delta\beta\| = 10$: SNR $\approx 10/0.47 \approx 21$. The regime change is detectable with high confidence.

2.3 Adaptive Threshold

The detection rule uses a z-score against rolling history:

$$z_t = \frac{d_t - \bar{d}_{t-H:t-1}}{s_{t-H:t-1}}$$

where $\bar{d}$ and s are the rolling mean and standard deviation of drift over the previous H windows. A regime change is flagged when $z_t > \tau$ (default $\tau = 2.5$).

The rolling history adapts to the current regime's volatility — a regime with higher natural fluctuation has a wider threshold.

2.4 Mode-Level Diagnosis

Beyond binary detection, the spectral approach identifies WHICH modes changed:

$$\delta_k = |A_{t,k} - A_{t-1,k}|$$

If mode 1 (largest eigenvalue direction) changed but modes 2–8 are stable: the dominant pattern shifted. If only high modes changed: fine structure altered, main pattern intact. This provides structural interpretation of the regime change.

---

3. Experiments

3.1 Synthetic Benchmark

Setup: $p = 10$ features, $T = 1200$ observations, 3 regimes:

Period	Regime	Coefficients
$t = 0$–500	A	$\beta_A = (5, 3, 0, \dots, 0)$
$t = 500$–800	B	$\beta_B = (-2, 0, 8, 0, \dots, 0)$ — different features active
$t = 800$–1200	A'	Back to β_A

Noise: $\epsilon \sim \mathcal{N}(0, 1.5^2)$. Window size: $n_w = 80$. Step: 40 samples. Shared basis: top 8 modes from the first window.

3.2 Detection Results

Window	Time	Drift	True regime	Detected?
0–10	0–480	0.3–1.0	A	stable
11	440–520	2.5	A→B transition	**DETECTED**
12	480–560	6.8	B	**DETECTED**
13–18	520–800	1.0–1.6	B	stable
19	760–840	5.9	B→A' transition	**DETECTED**
20	800–880	5.5	A'	residual
21–27	840–1160	0.4–0.8	A'	stable

True change points: $t = 500$, $t = 800$. **Detected at**: $t = 480$, $t = 760$. **Detection lag**: 20 and 40 samples (half a window to one window).

Within-regime drift: 0.3–1.0. **At transition**: 5.9–6.8. **Spike ratio: 6–10x.**

3.3 Drift Signature Analysis

The drift pattern has a characteristic shape at regime changes:

```
Regime A:  ~0.5   ~0.6   ~0.3   ~0.9   ~0.5    (normal fluctuation)
CHANGE:      2.5    6.8 ←                          two−step spike
Regime B:    1.1    1.0    1.2    1.6           (new normal, higher baseline)
CHANGE:      5.9    5.5 ←                          two−step spike
Regime A':     0.4    0.8    0.7                (back to original level)
```

The two-step spike occurs because the transition window straddles both regimes. The first window catches the edge; the second window is fully in the new regime.

3.4 Mode-Level Diagnosis

At the A→B transition ($t = 500$):

Mode	A_{before}	A_{after}	$\|\delta\|$	Interpretation
1	+5.1	-2.0	7.1	Dominant mode reversed sign
2	+3.0	+0.1	2.9	Second mode collapsed
3	+0.2	+7.9	7.7	Third mode activated
4–8	~ 0	~ 0	~ 0	Noise modes unchanged

The spectral diagnosis reveals: the regime change **activated a new feature** (mode 3) and **reversed the dominant pattern** (mode 1). This is structural information that no scalar change-point detector provides.

---

4. Comparison with Existing Methods

4.1 CUSUM

CUSUM monitors the cumulative sum of residuals from a fitted model. It detects mean shifts and linear trend changes but requires a pre-specified reference model. The spectral method requires no model specification — the eigenbasis IS the model.

4.2 Kernel Change-Point Detection (Harchaoui et al., 2009)

Kernel methods detect distributional changes in the feature space via MMD (Maximum Mean Discrepancy). They are powerful but provide no structural explanation of what changed. The spectral method decomposes the change into per-mode contributions.

4.3 Markov Switching (Hamilton, 1989)

Markov switching assumes a fixed number of regimes and estimates transition probabilities. It requires the number of regimes as input and assumes Markovian transitions. The spectral method is nonparametric — it detects any structural change without assuming a regime model.

4.4 Online PCA (Oja, 1982; Balsubramani et al., 2013)

Online PCA tracks the principal subspace over time. Related but different: online PCA detects changes in the data distribution $P(X)$, while spectral regime detection detects changes in the regression relationship $P(Y|X)$. The shared-basis approach preserves comparability that online PCA sacrifices.

---

5. Theoretical Analysis

5.1 Detection Power

Proposition 1 (Detectable regime change). A change in the coefficient vector $\Delta\beta = \beta_{new} - \beta_{old}$ is detectable (SNR > 1) if:

$$\|V_K^T \Delta\beta\| > \sigma\sqrt{K/n_w}$$

This gives the minimum detectable change as a function of noise, window size, and number of modes.

Corollary. For a given noise level, detection power improves with: (a) larger windows (n_w), (b) fewer modes (K), and (c) larger coefficient changes ($\|\Delta\beta\|$). The tradeoff: larger windows reduce false positives but increase detection lag.

5.2 False Positive Rate

Under the null hypothesis (no regime change), the drift d_t follows a χ distribution with K degrees of freedom, scaled by $\sigma/\sqrt{n_w}$. The z-score threshold $\tau = 2.5$ gives a false positive rate per window of approximately 0.6%. For 30 windows, the expected number of false positives is ~ 0.2.

5.3 Connection to Spectral Overfitting Theory

The regime detector uses K modes. The optimal K is determined by the spectral overfitting theory (Nagy, 2026c):

$$K^* = \Theta\left(\frac{\log(n_w/\sigma^2)}{\log\rho}\right)$$

Using $K > K^*$ adds noise modes to the drift statistic, increasing false positives without improving detection power. Using $K < K^*$ misses modes where the regime change occurred. The spectral overfitting diagnostic provides the optimal K for the detector.

---

6. Extensions

6.1 Multi-Scale Detection

Run detectors at multiple window sizes simultaneously: - Short windows ($n_w = 40$): fast detection, more false positives - Medium windows ($n_w = 100$): balanced - Long windows ($n_w = 250$): high confidence, delayed

A confirmed regime change is flagged at **all** time scales. This is equivalent to multi-resolution spectral analysis.

6.2 Kernel-Based Regime Detection

For nonlinear regime changes (e.g., the relationship becomes nonlinear after a structural break), use an RBF kernel eigenbasis instead of linear SVD. The method extends directly: the shared basis is the kernel eigendecomposition of the reference window.

6.3 Financial Applications

Portfolio risk: monitor spectral coefficients of the portfolio loss distribution. A regime change in the spectral state = the risk profile changed.

Trading strategies: if the spectral state of the return-generating process shifts, the strategy's edge may have disappeared. The drift statistic provides a quantitative signal for strategy rotation.

Credit risk: changes in the spectral structure of default correlations signal regime changes in credit markets (cf. 2008 correlation breakdown).

---

7. Conclusion

Spectral regime detection provides a nonparametric, interpretable method for identifying structural breaks by monitoring eigenmode coefficients over sliding windows. The method detects regime changes with 20–40 sample lag and 6–10x drift amplification, while providing per-mode diagnosis of what changed. The shared eigenbasis construction ensures coefficient comparability across windows, and the connection to the spectral overfitting theory provides the optimal number of monitoring modes.

The key advantage over existing methods is structural interpretability: not just "something changed" but "modes 1 and 3 jumped, mode 2 is stable." For financial applications, this translates directly to actionable information: which risk factors shifted, which remained stable, and by how much.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Bai, J. and Perron, P (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1), 47-78.
- Balsubramani, A., Dasgupta, S., and Freund, Y (2013). "The fast convergence of incremental PCA." *NeurIPS.* NeurIPS*.
- Hamilton, J. D (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2). DOI: 10.2307/1912559
- Harchaoui, Z., Moulines, E., and Bach, F (2009). Kernel change-point analysis. *NeurIPS.*.
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper.*
- Nagy, T. (2026). Spectral Model Compression: Provably Optimal Knowledge Extraction via Eigendecomposition. *Working paper.*
- Nagy, T. (2026). The Spectral Theory of Observation. *Working paper.*

- Oja, E (1982). A simplified neuron model as a principal component analyzer. *J. Mathematical Biology*, 15(3). DOI: 10.1007/bf00275687