

Constraint-Sweep Hysteresis in Rice Transcriptome State During Heat Stress and Recovery

A reproducible, theory-neutral test on GEO GSE74793

Jeremy C. Jones

HoldingLight LLC

ORCID: 0009-0007-2515-3774

contact@universalcollapse.com

Abstract

Biological responses to stress often exhibit path dependence: the trajectory during recovery may not retrace the trajectory during onset, even when the control parameter returns toward baseline. We operationalize this idea as a constraint-sweep hysteresis test on transcriptomic time series. Using GEO Series GSE74793 (rice leaf RNA-seq under heat and drought stress with recovery; Wilkins et al., 2016), we define a low-assumption “state” as PC1 of the submitter-provided processed expression matrix (DESeq2-normalized, variance-stabilized) and quantify hysteresis as the integral of the absolute separation between onset and recovery curves over their shared time window. A permutation null is constructed by shuffling phase labels within time bins, preserving temporal structure. In genotype-stratified analyses across four cultivars, heat stress exhibits strong hysteresis for Azucena and Pandan Wangi (permutation $p = 1/20,001 \approx 5.0 \times 10^{-5}$, permutation floor, at 20,000 permutations), while drought serves as a negative comparator and does not show comparable signal. All code, inputs, and outputs are packaged as a one-command reproducible outsider packet (a self-contained archive designed for independent replication without prior project context).

Keywords: hysteresis; path dependence; time series; transcriptomics; stress response; permutation test; reproducibility; rice

1. Introduction

Stress responses are not merely functions of present conditions; they carry memory of the path by which a system arrived there. In biology, such memory can be biochemical (metabolite pools), epigenetic (chromatin state), regulatory (network activity), or physiological (damage/repair dynamics). When memory exists, “recovery” may not simply reverse “onset.” This is the geometric signature of hysteresis: two trajectories occupy different regions of state space under the same value of a control parameter.

Here we implement a theory-neutral test for this signature in transcriptomic time series. We treat reported sample time as a practical constraint coordinate, compute a low-dimensional state from genome-wide expression, and quantify the loop between onset and recovery as a reproducible metric with a permutation-based p -value. Our focus is not on mechanistic attribution; it is on a robust, inspectable measurement of path dependence.

We evaluate the test on the rice (*Oryza sativa*) stress/recovery time series in GSE74793, originally generated as part of work on environmental gene regulatory influence networks (EGRINs; Wilkins et al., 2016).

2. Data

2.1 Dataset

We analyze GEO Series GSE74793, a transcriptome time series of 14-day rice leaves sampled every 15 minutes under heat stress and recovery, and drought stress and recovery, across multiple cultivars. The experimental design applies step-like perturbations: heat stress consists of a shift from 30°C to 40°C followed by return to 30°C; drought stress consists of air exposure followed by return to hydrated conditions (Wilkins et al., 2016). The associated BioProject (PRJNA301554) describes the full experimental design.

2.2 Acquisition and processing

Data were downloaded from NCBI GEO via the standard FTP directory structure. We use the Series matrix file for sample metadata (GSM accessions and characteristics) and the submitter-provided processed expression matrix (gene rows $\times$ sample columns). The processed matrix reflects DESeq2 median-of-ratios normalization and variance-stabilizing transformation (VST) as described in the original publication (Wilkins et al., 2016). We use this matrix directly, applying only gene-wise standardization (mean 0, variance 1 per gene across samples) before PCA. No additional log-transform is applied, as the VST already stabilizes variance on an approximately $\log_2$ -scale.

3. Methods

3.1 Overview

The test has five steps: (1) Build a runner metadata table labeling samples as onset (up) vs recovery (down) and assigning each sample a numeric constraint coordinate. (2) Compute a scalar state metric for each sample from the expression matrix. (3) Construct onset and recovery state-vs-constraint curves. (4) Quantify hysteresis via loop area plus auxiliary metrics. (5) Compute a permutation p-value using a time-bin phase shuffle.

All steps are implemented in `analysis_gene_expression_hysteresis.py` and companion scripts included in the outsider packet.

3.2 Runner metadata construction

We parse `Sample_characteristics_ch1` fields in the GEO metadata and extract treatment labels (HEAT, RECOV_HEAT, DROUGHT, RECOV_DROUGHT, CONTROL) and sample time in minutes. We then define two analyses: heat onset vs heat recovery, and drought onset vs drought recovery.

For each analysis, the **constraint coordinate** is the reported sample time (minutes), and **phase** encodes which constraint-regime the system occupies (stress vs recovery). The constraint coordinate is not the physical stressor itself (temperature, water potential) but the temporal position within the experimental protocol. This is a practical choice: the stressor transitions are step-like rather than continuously graded, so time serves as the natural axis along which state evolution can be tracked and compared between phases.

Constraint overlap is required. The method only compares curves over the shared time window where both phases have defined values. In this dataset, the heat overlap window spans 135–240 minutes (105 min range) and the drought overlap window spans 105–135 minutes (30 min range).

3.3 State metric

Let X be the submitter-provided processed expression matrix (DESeq2-normalized + VST; genes $\times$ samples). We standardize each gene across samples to mean 0 and variance 1, then compute PCA with two components. The state metric is the PC1 score per sample. In genotype-stratified analyses across all four cultivars (Azucena, Pandan Wangi, Tadukan, and Kinandang Puti), PC1 explains 14.1–15.4% of total variance; in the pooled heat analysis, 9.6%. These values are typical for high-dimensional transcriptomic data where thousands of genes contribute independent variation. This choice is intentionally minimal: it produces a single scalar state per sample without assuming a specific gene set or pathway model. That the hysteresis signal emerges from a metric capturing only $\sim 14\%$ of total variance suggests the path-dependence effect is strong enough to dominate the leading axis of variation.

3.4 Curve construction

Within each phase, samples are sorted by constraint coordinate, duplicate time values are collapsed by averaging state, and both phase curves are interpolated linearly onto a common grid of 200 evenly spaced points over the overlap window.

3.5 Hysteresis metrics

Let $s_u(c)$ be the onset curve and $s_d(c)$ the recovery curve over the overlap domain $[c_{\min}, c_{\max}]$. We define:

Loop area: $A = \int |s_u(c) - s_d(c)| dc$, computed by trapezoidal integration on the interpolation grid.

Normalized loop area: $A_{\text{norm}} = A / (c_{\max} - c_{\min})$, enabling comparison across analyses with different overlap geometries (heat: 105 min vs drought: 30 min).

Lag index: Mean absolute separation on the interpolation grid: $L = (1/N) \sum |s_u(c_i) - s_d(c_i)|$. Lag index is computed directly on the discrete grid and is the native output of the analysis code. $A_{\text{norm}} = A / (c_{\max} - c_{\min})$ is the continuous-normalized equivalent. Both are reported: lag index for reproducibility against the code output, A_{norm} for cross-condition comparability.

Return distance: Absolute difference between early baseline state (earliest ~10% of samples) and late recovery state (latest ~10% of down-phase samples).

z-score: $z = (A_{\text{obs}} - \mu_{\text{null}}) / \sigma_{\text{null}}$, where μ_{null} and σ_{null} are the mean and standard deviation of the permutation null distribution. Provides a scale-independent measure of how far the observed area departs from the null.

3.6 Permutation null and p-value

To preserve the temporal structure, we generate null areas by shuffling phase labels within each time bin over the overlap window: for each unique time value c within $[c_{\min}, c_{\max}]$, permute the phase label (up/down) among all samples sharing that time value. Bins with ≤ 1 sample are skipped. Loop area A is then recomputed on the shuffled metadata. This is repeated for n_{perm} seeded permutations.

The p-value is: $p = (1 + \#\{A_{\text{null}} \geq A_{\text{obs}}\}) / (1 + n_{\text{perm}})$. When no null value reaches or exceeds the observed area, the plus-one estimator yields $p = 1/(n_{\text{perm}} + 1)$, the resolution floor of the test. At 20,000 permutations, this floor is $p = 1/20,001 \approx 5.0 \times 10^{-5}$. Throughout this paper, we report this value as the permutation floor whenever zero null values exceed A_{obs} , to distinguish floor-hitting results from exact p-values.

3.7 Two-stage genotype-stratified analysis

We parse genotype/cultivar labels from sample characteristics and stratify into four subsets: Azucena, Pandan Wangi, Kinandang Puti, and Tadukan. The analysis proceeds in two pre-specified stages:

Stage 1 (primary screen): 8-test screen (4 genotypes $\times$ 2 conditions) at 2,000 permutations, with Benjamini–Hochberg (BH) FDR correction across all 8 tests. Any condition with at least one genotype reaching $q < 0.05$ advances its full genotype family to Stage 2 for within-family comparability.

Stage 2 (tail refinement): The advancing family is re-run at 20,000 permutations to tighten tail probability estimates. BH correction is applied within the refinement family. Conditions with no flagged genotypes remain at Stage 1. This two-stage design is declared in advance and controls the false discovery rate within each stage via BH correction.

4. Results

4.1 Whole-series (pooled) hysteresis

Pooled over genotypes, both analyses execute cleanly with non-empty overlap windows. Heat (pooled) yields $A \approx 2,178$, $z = 0.49$, $p \approx 0.309$ (2,000 permutations; $\mu_{\text{null}} = 1,940$, $\sigma_{\text{null}} = 490$). Drought (pooled) yields $A \approx 721$, $z = 0.14$, $p \approx 0.371$ (2,000 permutations; $\mu_{\text{null}} = 691$, $\sigma_{\text{null}} = 210$). The pooled p-values are not extreme, consistent with signal heterogeneity across genotypes. Both pooled z-scores are below 0.5, confirming that genotype aggregation dilutes the signal.

4.2 Stage 1: primary screen (2,000 permutations)

The 8-test screen identifies two heat genotypes at the permutation floor ($p = 1/2,001 \approx 5.0 \times 10^{-4}$, $q \approx 0.002$ across 8 tests): Azucena and Pandan Wangi. All drought tests and remaining heat tests yield $q \geq 0.20$. Per protocol, because heat contains at least one flagged genotype, the full heat family (all 4 genotypes) advances to Stage 2 for within-family comparability. Drought remains at Stage 1.

4.3 Stage 2: tail refinement (heat; 20,000 permutations)

Table 1 reports the full results of the 20,000-permutation heat refinement. Two genotypes exhibit extreme hysteresis signals: Azucena ($A = 5,316$; $p = 5.0 \times 10^{-5}$, permutation floor) and Pandan Wangi ($A = 4,934$; $p = 5.0 \times 10^{-5}$, permutation floor). In both cases, zero of 20,000 null permutations reached or exceeded the observed loop area. Kinandang Puti and Tadukan show moderate, non-significant effects. BH correction within the heat family of 4 tests yields $q \approx 1.0 \times 10^{-4}$ for both significant genotypes.

Table 1. Stage 2 heat hysteresis results (20,000 permutations). BH-corrected within heat family (4 tests).

Genotype	Area	A_{norm}	p	q (BH)	z	Lag Idx	Ret. D.	n $\uparrow/\downarrow$
Azucena	5,316	50.63	$5.0 \times 10^{-5}\dagger$	$\approx 1.0 \times 10^{-4}$	4.79	50.58	229.8	32/16
Pandan Wangi	4,934	46.99	$5.0 \times 10^{-5}\dagger$	$\approx 1.0 \times 10^{-4}$	4.57	46.96	246.3	32/16
Tadukan	5,063	48.22	0.102	0.102	1.33	48.29	154.1	32/16
Kinandang Puti	3,617	34.45	0.084	0.102	1.43	34.43	231.9	32/15

Note: $\dagger$ Permutation floor: $0/20,000$ null $\geq \text{obs}$; $p = 1/(n_{\text{perm}} + 1) = 1/20,001$. $A_{\text{norm}} = A / 105$ min. All p-values and z-scores computed from the same Stage 2 (20,000-perm) null distributions. BH within heat family (4 tests). $z = (A_{\text{obs}} - \mu_{\text{null}}) / \sigma_{\text{null}}$.

4.4 The Tadukan puzzle

An instructive feature of the data is that Tadukan exhibits a large loop area ($A = 5,063$; $A_{\text{norm}} = 48.22$) numerically comparable to Azucena ($A = 5,316$; $A_{\text{norm}} = 50.63$), yet its permutation p-value (0.102) is non-significant. The z-scores reveal why: Azucena's $z = 4.79$ ($\mu_{\text{null}} = 1,936$, $\sigma_{\text{null}} = 706$) while Tadukan's $z = 1.33$ ($\mu_{\text{null}} = 3,952$, $\sigma_{\text{null}} = 837$). Tadukan's null distribution has more than double the mean and a wider spread—the permutation shuffle generates much larger expected loop areas under the null for this genotype's particular sample structure. Despite similar raw areas, Azucena sits nearly 5 standard deviations above its null while Tadukan sits barely above 1. The permutation framework correctly calibrates each genotype's result against its own null geometry.

4.5 Drought

All four genotypes under drought stress yield non-significant hysteresis at 2,000 permutations ($p \approx 0.28$ – 0.92 ; $q \geq 0.44$ across the 8-test screen). The pooled drought z-score ($z = 0.14$) confirms the observed loop area is indistinguishable from the null mean. The drought overlap window is substantially narrower (30 min vs 105 min for heat), limiting the domain over which hysteresis can accumulate. Whether this null result reflects genuinely smaller path dependence, a mismatch between the informative biological window and the measured overlap, or a limitation of the PC1 state metric under drought conditions remains a testable question for future work.

Figures




Figure 1. Constraint-sweep hysteresis for Azucena under heat stress. State (PC1 of gene-standardized VST expression) vs time (minutes) for onset (blue) and recovery (orange) over the shared overlap window (135–240 min). The persistent separation between curves indicates path-dependent transcriptomic state. $A = 5,316$; $p = 5.0 \times 10^{-5}$, permutation floor (20,000 perms).




Figure 2. Permutation null distribution of loop areas for Azucena under heat stress (20,000 permutations, time-bin phase shuffle). Vertical line marks the observed loop area. No null value reaches the observed area ($p = 1/20,001$, permutation floor). The null distribution centers near $\mu \approx 1,936$ ($\sigma \approx 706$) while the observed area exceeds 5,300 ($z = 4.79$).




Figure 3. Constraint-sweep hysteresis for Pandan Wangi under heat stress. Same analysis as Figure 1. Independent genotype replication of the path-dependence signal. $A = 4,934$; $p = 5.0 \times 10^{-5}$, permutation floor (20,000 perms).




Figure 4. Permutation null distribution for Pandan Wangi under heat stress (20,000 permutations). Observed loop area sits in the extreme right tail with no null value reaching it. Confirms the heat hysteresis signal replicates across genotypes.




Figure 5. Constraint-sweep hysteresis for drought stress (pooled across genotypes). State vs time over the narrow overlap window (105–135 min). Onset curve diverges sharply after ~120 min while recovery remains flat. Despite visual separation, the narrow overlap domain (30 min) and small sample sizes yield a non-significant result ($p \approx 0.371$, $z = 0.14$).




Figure 6. Permutation null distribution for drought stress (pooled, 2,000 permutations). Observed loop area (vertical line) sits within the body of the null distribution ($\mu_{\text{null}} = 691$, $\sigma_{\text{null}} = 210$). Inclusion of this negative result demonstrates that the method discriminates rather than producing universally positive signals.

5. Discussion

5.1 What the test detects

This test detects path dependence in transcriptomic state under a shared temporal coordinate. In heat stress, two of four genotypes show robust separation between onset and recovery trajectories, producing loop areas incompatible with a time-bin-preserving phase shuffle null. The system does not retrace its trajectory during recovery—it retains path dependence throughout.

5.2 Why pooling can wash out signal

The pooled analysis aggregates genotypes with heterogeneous dynamics. A mixture of strong and weak hysteresis responses yields a moderate pooled p-value (0.309, $z = 0.49$) even when a subset exhibits extreme effects. The genotype-stratified analysis resolves this heterogeneity, revealing that the heat signal is concentrated in Azucena and Pandan Wangi.

5.3 Overlap geometry and cross-condition comparability

The heat and drought analyses differ substantially in overlap geometry: 105 minutes for heat vs 30 minutes for drought. Raw loop area is therefore not directly comparable across conditions. The normalized metric $A_{\text{norm}} = A / (c_{\text{max}} - c_{\text{min}})$ addresses this by expressing hysteresis as mean state-separation per unit time. Even after normalization, heat produces larger effects than drought, but the narrower drought window limits statistical power and should be considered when interpreting the null result.

5.4 Limitations and sensitivity checks

State metric choice: PC1 is convenient but not unique; alternative state constructions (gene sets, diffusion maps, supervised contrasts) may yield different sensitivity profiles. **Constraint coordinate:** we use reported sample time rather than a continuous physical measurement of the stressor (temperature, water potential). Designs with continuously monitored stressor values would enable a true constraint axis rather than a temporal proxy. **Permutation design:** we shuffle phases within time bins over the overlap window; alternative nulls (block permutations, genotype-condition constrained shuffles) can be evaluated. **Overlap dependence:** hysteresis is only measured on the shared time window; explicit experimental designs with symmetric sweeps would increase power.

5.5 Broader framing

Although this paper is intentionally theory-neutral, the same measurable signature—constraint-conditioned hysteresis—can serve as a cross-domain marker of structured memory under constraint. In the broader Universal Collapse Theory (UCT) program, we treat this as a candidate empirical signature of constraint-conditioned stabilization (designated S_3 : sweeps $\rightarrow$ hysteresis), but the statistical test stands independently of that interpretation. The method is portable to any domain with time-series data under sweepable constraints.

6. Reproducibility and Availability

GEO accession: GSE74793. **BioProject:** PRJNA301554.

All code, scripts, and outputs are packaged as a reproducible outsider packet, archived under the Empirical Demonstrations component of the UCT OSF project (DOI: <https://doi.org/10.17605/OSF.IO/KZ8TP>). The Makefile workflow provides the following targets:

```
make download  # fetch GEO data from NCBI
```

```
make run  # execute full analysis pipeline
```

```
make figures  # generate all figures and tables
```

Primary artifacts include per-genotype results CSVs, metadata JSONs, hysteresis curve figures, and null distribution figures, organized by condition and genotype. Fixed random seeds ensure exact reproducibility.

References

Wilkins, O., Bračun, S., Naseem, H., et al. (2016). EGRINs (Environmental Gene Regulatory Influence Networks) in Rice That Function in the Response to Water Deficit, High Temperature, and Agricultural Environments. *The Plant Cell*, 28(10), 2365–2384. doi:10.1105/tpc.16.00158

NCBI BioProject PRJNA301554.

NCBI GEO Series GSE74793.

Appendix A: Full Numerical Results

Table A1. Complete heat results (Stage 2, 20,000 permutations).

Genotype	Area (obs)	A _{norm}	p-value	q (BH)	Lag Idx	Ret. Dist.	Samples
Azucena	5,316.29	50.63	5.00×10⁻⁵†	≈1.0×10⁻⁴	50.581	229.833	48 (32↑/16↓)
Pandan Wangi	4,934.44	46.99	5.00×10⁻⁵†	≈1.0×10⁻⁴	46.964	246.252	48 (32↑/16↓)
Tadukan	5,063.46	48.22	0.10195	0.10195	48.292	154.097	48 (32↑/16↓)
Kinandang Puti	3,616.59	34.45	0.08380	0.10195	34.431	231.900	47 (32↑/15↓)

All: grid_n = 200; seed = 7; n_{perm} = 20,000. Overlap: 135–240 min (105 min). BH within heat family (4 tests).
Expression: DESeq2 + VST (submitter-provided). State: PC1 of gene-standardized expression. †Permutation floor: 0/20,000 null ≥ obs; p = 1/(n_{perm} + 1) = 1/20,001.

Appendix B: Pooled Diagnostics

Table B1. Pooled (all-genotype) results with null distribution diagnostics (2,000 permutations).

Condition	A _{obs}	p	z	μ _{null}	σ _{null}	Interpretation
Heat	2,178	0.309	0.49	1,940	490	<0.5σ above null
Drought	721	0.371	0.14	691	210	Indistinguishable

The pooled heat z = 0.49 reflects signal heterogeneity: strong hysteresis in Azucena and Pandan Wangi is diluted by weaker effects in Tadukan and Kinandang Puti. Genotype stratification resolves this.

Appendix C: Analysis Parameters

For exact reproducibility, the following parameters are fixed in the analysis configuration:

Expression source: GEO-provided processed matrix (DESeq2 + VST; Wilkins et al., 2016)

Gene standardization: per-gene mean = 0, variance = 1 across all samples in the analysis subset

Dimensionality reduction: PCA (2 components); state = PC1 score

Interpolation grid: 200 evenly spaced points over the overlap window

Permutation null: time-bin phase shuffle within overlap window; bins with ≤1 sample skipped

Random seed: 7 (fixed for all permutation runs)

Stage 1: n_{perm} = 2,000; BH correction across 8 tests (4 genotypes × 2 conditions)

Stage 2: $n_{\text{perm}} = 20,000$; BH correction within advancing condition family (heat: 4 tests)

AI Disclosure. AI tools were used to assist with manuscript preparation. The underlying theory, arguments, and interpretive claims are the author's own, and the author takes full responsibility for the content.

Citation: Jones, J. C. (2026). Constraint-Sweep Hysteresis in Rice Transcriptome State During Heat Stress and Recovery: A Reproducible, Theory-Neutral Test on GEO GSE74793. HoldingLight LLC.

This paper is part of the Universal Collapse Theory library.

For a reading guide and full architecture, visit universalcollapse.com/roadmap