Galaxy Structural Fingerprints: A Multiscale Atlas of Persistence, Fragmentation, and Connectivity

UNNS Substrate Research Collective[1]

[1]UNNS Structural Astrophysics Programme

August 3, 2026

Abstract

Rotation curves are normally compared through fitted mass models or empirical acceleration relations; the multiscale structural identity of the underlying baryonic-acceleration profile is rarely studied as an object in its own right. We introduce the UNNS structural-response fingerprint and use it to establish structural fingerprints of galaxy baryonic-acceleration ladders derived from resolved SPARC mass models: for each galaxy, the valid observed radial measurements are used to compute $g_{\mathrm{bar}}(R)$, and the chamber input is the numerically ascending vector of unique $\log_{10} g_{\mathrm{bar}}$ values – the baryonic-acceleration ladder. The resulting multichannel object $\mathcal{F}_g(\kappa) = \{\rho_g(\kappa),\, D_{A,g}(\kappa),\, \tilde{\nu}_g(\kappa)\}$, with $D_{A,g}(\kappa) = 1 - A_{\kappa,g}(\kappa)$, is computed across a logarithmically spaced perturbation scale κ, and we construct a fingerprint atlas for the 121-galaxy primary SPARC sample. We report near-perfect unchanged-preparation retrieval (top-1 and top-5 source recovery of 1.000, median intraclass correlation 0.99985), a channel-dependent low-dimensional architecture (ρ and ν retain 90.8% and 88.5% of variance in three components; the more heterogeneous A_κ deficit retains 55.9%), a discrete Weak Persistence regime separated from the main Stable Structure population, a five-family consensus admissibility structure within the Stable Structure basin, with strong internal co-assignment coherence, moderate-to-good resampling stability, and significant differentiation in independent physical and connectivity properties, physical neighbour coherence relative to an all-pair baseline, and an empirical separation of persistence, fragmentation, and connectivity (110/121 galaxies are structurally stable, 103/121 show hard fragmentation, and only 17/121 reach full percolation, exclusively through an adaptive κ extension). We further report the fingerprint's measurable response to observational representation (Gate C) and an external transfer test against the Radial Acceleration Relation and its local baryonic-acceleration baselines, in which V_{obs} enters only later, through g_{obs}, as the external transfer target. Static whole-galaxy chamber signatures do not improve held-out radial prediction beyond a simple local baseline, while the full $\rho(\kappa)$ curve emerges as the leading transfer-sensitive channel, improving prediction in four of five folds without clearing the full registered confirmation gate. We conclude that galaxy baryonic-acceleration ladders possess reproducible, organized, multichannel structural identities, and that using this identity to predict gravitational discrepancy requires local, regime-conditioned, and channel-selective architectures rather than a single static whole-galaxy descriptor.

Keywords: galaxy structural fingerprint, multiscale response morphology, structural-fingerprint atlas, structural identity, persistence–fragmentation–connectivity, Radial Acceleration Relation, SPARC

1 Introduction

1.1 Galaxy rotation curves and baryon–dynamics coupling

A galaxy rotation curve records how the observed circular velocity $V_{\rm obs}(R)$ varies with galactocentric radius R. Decades of resolved H I and Hα kinematics, combined with near-infrared photometry, have shown that this curve is not an arbitrary function: it is tightly coupled to the distribution of a galaxy's baryonic mass, and the shape of that coupling is itself a rich source of structural information that conventional analysis rarely interrogates directly.

Conventional analysis of a rotation curve asks a narrow set of questions: what mass distribution fits the curve; what empirical relation predicts the observed acceleration from the baryonic acceleration; and what morphological class the galaxy belongs to. These questions compress the curve into a small number of fitted or classified quantities. This paper asks a different, complementary question: *what multiscale structural identity does the complete rotation curve express, independent of any specific fitted law?* Two galaxies with visually similar rotation curves may nonetheless respond differently once their sampled baryonic-acceleration values are interrogated under scale-dependent structural perturbation, and that differential response is itself a measurable, reproducible property of the galaxy.

1.2 The Radial Acceleration Relation as the empirical baseline

Before introducing the structural-fingerprint programme, we fix the strongest existing empirical regularity in galaxy dynamics, since it supplies both context and, later, the null hypothesis against which fingerprint-based prediction is tested. At each radius R, define the observed and baryonic centripetal accelerations

$$g_{\rm obs}(R) = \frac{V_{\rm obs}^2(R)}{R}, \qquad g_{\rm bar}(R) = \frac{V_{\rm bar}^2(R)}{R}. \tag{1}$$

Across rotationally supported galaxies spanning wide ranges of mass, morphology, size, surface brightness, and gas fraction, $g_{\rm obs}$ is tightly organized by $g_{\rm bar}$: this is the Radial Acceleration Relation (RAR). The original SPARC analysis demonstrated the relation using 2,693 radial measurements across 153 galaxies (McGaugh et al., 2016; Lelli et al., 2016a); the complete SPARC database comprises 175 disk galaxies with Spitzer 3.6 μm photometry and resolved H I/Hα rotation curves (Lelli et al., 2016b). A widely used empirical representation is

$$g_{\rm obs} = \frac{g_{\rm bar}}{1 - \exp\left[-\sqrt{g_{\rm bar}/g_\dagger}\right]}, \qquad g_\dagger \approx 1.2 \times 10^{-10}\ {\rm m\,s^{-2}}, \tag{2}$$

which reduces to $g_{\rm obs} \simeq g_{\rm bar}$ at high acceleration and to the deep-regime scaling $g_{\rm obs} \simeq \sqrt{g_{\rm bar}\, g_\dagger}$ at low acceleration. Reported scatter is below approximately 0.13 dex in the original ensemble analysis, and closer to 0.057 dex root-mean-square in later individual-galaxy fits that marginalize over mass-to-light ratio, distance, and inclination (Li et al., 2018). The characteristic scale $g_\dagger$ is consistent in magnitude with the characteristic Modified Newtonian Dynamics (MOND) acceleration a_0 (Milgrom, 1983); the values are closely related, but their precise numerical identification depends on the adopted interpolation function and modelling assumptions, and MOND does not uniquely select the specific empirical exponential interpolation used above. Alternative accounts interpret the RAR as an emergent regularity of coupled baryonic and dark-matter halo assembly within ΛCDM (Di Cintio & Lelli, 2016; Navarro et al., 2017), and recent joint kinematic and weak-lensing analysis extends the RAR by approximately 2.5 dex toward lower accelerations and finds that early- and late-type galaxies occupy the same joint relation when strict isolation criteria and consistent stellar- and gas-mass estimates are applied (Mistele et al., 2024). The origin and precise regularity of the RAR remain active subjects of study; we take no position on that debate here.

1.3 What the RAR does not describe: whole-ladder structural identity

The RAR relates local $g_{\mathrm{bar}}(R)$ to local $g_{\mathrm{obs}}(R)$ at matched radius. The chamber addresses a different object: the global spacing geometry of the sampled baryonic-acceleration ladder after ascending canonicalization. It tests whether the complete set of sampled baryonic accelerations possesses reproducible multiscale response morphology, structural regimes, and connectivity properties. This does not weaken the discovery; it identifies the discovery precisely: **the manuscript establishes structural fingerprints of galaxy baryonic-acceleration ladders derived from resolved SPARC mass models.** These are structural questions about the whole ladder, not local dynamical questions about matched-radius acceleration, and they are the subject of this paper.

1.4 UNNS multiscale structural fingerprints

We represent each galaxy g by a multichannel response object measured across a chamber scale κ,

$$\mathcal{F}_g(\kappa) = \left[\, \rho_g(\kappa),\ D_{A,g}(\kappa),\ \tilde{\nu}_g(\kappa) \,\right], \qquad D_{A,g}(\kappa) = 1 - A_{\kappa,g}(\kappa), \tag{3}$$

constructed from the sorted, deduplicated baryonic-acceleration ladder of the galaxy under a battery of controlled perturbation and admissibility tests (Section 2). Here $D_{A,g}(\kappa)$ is the admissibility deficit and $\tilde{\nu}_g(\kappa) = \nu_g(V_\kappa;\kappa)/\max(\lceil (n_g-1)/2 \rceil,\, 1)$ is the ladder-size-normalized inversion budget used in the registered combined-distance metric, where n_g is the number of unique ladder values for galaxy g; for intuitive readability, Fig. 11 instead plots the raw forms $A_\kappa(\kappa)$ and $\nu(V_\kappa)(\kappa)$, which carry equivalent information under a monotonic rescaling. Equation (3) is not one verdict, one PCA coordinate, or one fitted number: it is the complete multiscale response morphology produced as κ is swept over its full range. Scalar summaries, PCA coordinates, and chamber verdicts each describe a projection of this object; the full curves preserve its richest content. We refer to $\mathcal{F}_g(\kappa)$ throughout as the galaxy's *structural fingerprint*, and to the collection of fingerprints, together with the coordinate space and neighbour relation built from them, as the *structural-fingerprint atlas*.

1.5 The two independent questions

The programme separates two questions that are easy to conflate:

(i) **Structural identity.** Do galaxies possess reproducible, galaxy-specific multiscale structural fingerprints, organized into an interpretable atlas?

(ii) **Dynamical transfer.** Do those fingerprints add held-out predictive information for gravitational discrepancy beyond RAR-like local baselines?

These questions do not have to receive the same answer, and, as we show below, they do not: the fingerprints are established as reproducible and organized (question i), while a single static whole-galaxy representation does not demonstrate incremental transfer beyond the tested local baselines (question ii). The two results are complementary rather than contradictory, and together they define both the discovery and its boundary.

1.6 Main contributions

1. Construction of a 121-galaxy structural-fingerprint atlas from the primary SPARC sample.
2. Demonstration of galaxy-specific response identity under unchanged chamber repetition.
3. Discovery of channel-specific atlas architecture: ρ and ν are strongly compressible, A_κ is more heterogeneous.
4. Separation of persistence, fragmentation, and connectivity into three empirically distinct structural dimensions.
5. Identification of a five-family consensus admissibility structure within the Stable Structure basin, with strong internal co-assignment coherence, moderate-to-good resampling stability, and statistically significant differentiation in independent physical and connectivity properties.
6. Demonstration of physical alignment and cross-class structural analogues in the atlas.
7. Mapping of the observational transformation layer under radius deletion, thinning, coverage change, and mass-to-light preparation.
8. Identification of an external transfer boundary and of $\rho(\kappa)$ as the leading transfer-sensitive channel for future local, regime-conditioned architectures.

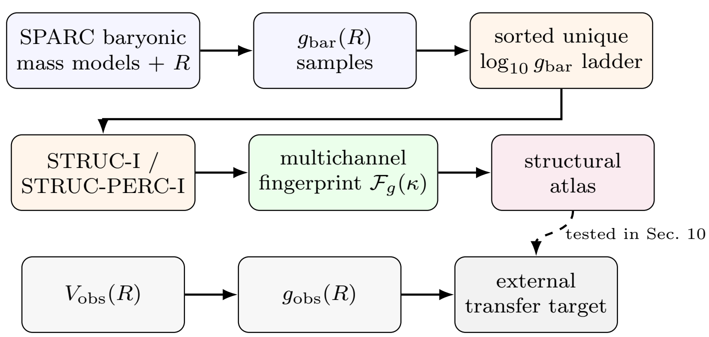

Figure 1: Pipeline from SPARC baryonic mass models to the structural-fingerprint atlas. Each galaxy's valid radial measurements are converted to $g_{\rm bar}(R)$ and canonicalized into a sorted, deduplicated $\log_{10} g_{\rm bar}$ ladder, which is passed through the STRUC-I and STRUC-PERC-I chambers across a logarithmic κ grid to produce the multichannel fingerprint $\mathcal{F}_g(\kappa)$ and the structural atlas. The observed rotation curve enters only through the independent $g_{\rm obs}(R)$ branch, which supplies the external transfer target tested against the atlas in Section 10; it is not part of the structural fingerprint.

Figure 1 summarizes this pipeline. Section 2 defines the sample and the chamber constructions; Section 3–Section 8 report the atlas itself; Section 9 reports its observational transformation layer; Section 10 reports the external transfer test and the RAR-based null; Section 11 discusses the results and Section 12 concludes.

Table 1: Primary galaxy sample and input-variable summary ($n = 121$).

Quantity	Median	Minimum	Maximum
Distance (Mpc)	18.0	1.0	127.8
Inclination (°)	64	15	90
3.6 μm luminosity ($10^9\,L_\odot$)	10.8	0.01	490.0
Effective surface brightness ($L_\odot\,\mathrm{pc}^{-2}$)	157.3	6.9	2709.1
Flat rotation velocity ($\mathrm{km\,s^{-1}}$)	113.0	0	332.0
H I mass ($10^9\,M_\odot$)	2.68	0.07	40.1
Valid radial points	19	10	115

A flat rotation velocity of 0 $\mathrm{km\,s^{-1}}$ follows the SPARC catalogue convention for rotation curves without a measured flat portion; it is a missing-value flag, not a physical zero velocity.

2 Data and Structural-Fingerprint Construction

2.1 SPARC galaxy sample

The primary sample comprises 121 galaxies drawn from the Spitzer Photometry and Accurate Rotation Curves (SPARC) database (Lelli et al., 2016b), selected under a frozen rule: SPARC quality flag $Q \leq 2$ and at least ten valid observed radial points. Table 1 summarizes the resulting sample. The sample spans four orders of magnitude in 3.6 μm luminosity, distances from below 1 to nearly 130 Mpc, inclinations from 15° to 90°, and Hubble types from 0 to 11, i.e. from early-type disks to irregulars. No galaxy was added to or removed from the primary sample after protocol registration.

2.2 Baryonic-acceleration ladder representation

For each galaxy, the valid observed radial measurements are used to calculate $g_{\mathrm{bar}}(R)$. The chamber input is then constructed as the numerically ascending vector of unique $\log_{10} g_{\mathrm{bar}}$ values,

$$\mathrm{ladder}_g = \mathrm{unique}(\mathrm{sort}(\log_{10} g_{\mathrm{bar}}))\,. \tag{4}$$

Thus the ladder preserves the complete sampled baryonic-acceleration value set and its adjacent-gap geometry, but not the original radial ordering: radial order is not retained in the chamber input, and duplicate acceleration values are removed. V_{obs} is not part of the structural fingerprint; it enters later, through g_{obs}, only as the external transfer target of Section 10. No interpolation or extrapolation is used. For deletion, thinning, and coverage experiments (Section 9), radial positions are first selected according to the registered branch rule, after which the ladder is rebuilt using the same ascending sort-and-unique operation.

2.3 STRUC-I chamber

STRUC-I v1.0.4 measures the perturbative admissibility of the canonical (sorted, deduplicated) ladder across a κ sweep of 40 logarithmically spaced points from $\kappa_{\mathrm{min}} = 0.01$ to $\kappa_{\mathrm{max}} = 1.0$. At each κ, the ladder gaps are perturbed by $\delta_i \sim \mathrm{Uniform}[-\varepsilon, \varepsilon]$ with $\varepsilon = \kappa \times \mathrm{median}(\text{positive adjacent gaps})$, using 2,000 Monte Carlo perturbation draws per κ point. From this sweep the chamber exports three channels: $A_\kappa(\kappa)$, the fraction of perturbation draws whose inversion count remains within the admissible vulnerability budget at scale κ; $\nu(V_\kappa)(\kappa)$, the vulnerability-derived inversion budget itself; and $\rho(\kappa)$, the chamber's primary structural-response coordinate, tracking the degree of structural response at scale κ (full computational definitions are given in the STRUC-I protocol; UNNS Structural Programme 2026). Persistence classification (Stable Structure versus Weak Persistence) is derived from the shape of this three-channel response across the full κ sweep.

2.4 STRUC-PERC-I chamber

STRUC-PERC-I v2.5.0 measures full-pairwise gap connectivity and fragmentation on the same canonical ladder, independent of STRUC-I admissibility. It reports a fragmentation verdict (e.g. HARD_FRAGMENTATION, FULL_PERCOLATION, TAIL_FRAGMENTATION), the connectivity scope reached on the native κ range, and, where the native range is insufficient, the adaptive extension factor κ-connect required to reach full connectivity. Admissibility and connectivity are computed independently and are not assumed to agree.

Connectivity and fragmentation refer to the graph formed from similarities among adjacent-gap magnitudes in the ordered baryonic-acceleration ladder. They do not refer to spatial connectivity, morphological fragmentation, or fragmentation of the galactic disk. This distinction matters because the manuscript title elevates fragmentation and connectivity to central concepts, and both terms could otherwise be misread as describing the galaxy's physical structure rather than the response geometry of its sampled acceleration values.

2.5 Fingerprint representation

Each galaxy's fingerprint is available in four representations of decreasing richness and increasing compactness: (i) the complete 40-point three-channel curve $\mathcal{F}_g(\kappa)$; (ii) compact per-galaxy summary statistics (e.g. mean, maximum, and value-at-$\kappa_{\max}$ for each channel); (iii) PCA coordinates computed separately per channel and on the combined three-channel signature; and (iv) pairwise Euclidean distances and nearest-neighbour relations computed in PCA space. Table 2 collects formal definitions used throughout.

2.6 Validation gates

Three validation stages structure the results below. **Gate B** tests whether unchanged-preparation repetition of STRUC-I returns the same galaxy-specific fingerprint far more closely than the fingerprint of any other galaxy (Section 3). **Gate C** tests how the fingerprint transforms under four controlled changes to observational representation: single-radius deletion, deterministic thinning, fixed-count coverage change, and mass-to-light preparation sensitivity (Section 9). The **external transfer validation** tests whether static or full-curve chamber signatures add held-out predictive information for gravitational discrepancy beyond RAR-like local baselines (Section 10).

Table 2: Definitions of fingerprint channels and structural metrics.

Symbol / term	Definition
κ	Chamber perturbation scale; 40 points log-spaced on [0.01, 1.0].
$\rho(\kappa)$	Primary STRUC-I structural-response coordinate at scale κ.
$\nu(V_\kappa)(\kappa)$	Inversion-vulnerability budget at scale κ, derived from V_κ.
$A_\kappa(\kappa)$	Fraction of 2,000 Monte Carlo perturbation draws at scale κ whose inversion count remains admissible.
$D_{A,g}(\kappa)$	Admissibility deficit, $D_{A,g}(\kappa) = 1 - A_{\kappa,g}(\kappa)$; the fingerprint channel used in Eq. (3) in place of $A_\kappa(\kappa)$.
$\tilde{\nu}_g(\kappa)$	Ladder-size-normalized inversion budget, $\nu_g(V_\kappa;\kappa)/\max(\lceil(n_g - 1)/2\rceil, 1)$.
$\mathcal{F}_g(\kappa)$	Complete multichannel fingerprint, Eq. (3).
Structural persistence P_g	Stability of the multiscale response across the κ sweep (Stable Structure vs. Weak Persistence).
Fragmentation Φ_g	STRUC-PERC-I verdict on gap connectivity (e.g. Hard Fragmentation).
Connectivity reach C_g	Whether, and at what κ-connect extension, full percolation is reached.
Structural-PCn	nth principal component of the combined three-channel PCA.
ρ-PCn	nth principal component of the $\rho(\kappa)$ channel alone.

3 A Reproducible Galaxy Structural-Fingerprint Atlas

3.1 Galaxy-specific source recovery

We first establish that the fingerprint exists as a reproducible object before analysing what it means. Five unchanged-ladder STRUC-I executions (R1–R5) were frozen under a registered protocol using the chamber's native, unseeded random-number generator; independence across repeats is represented by execution timestamp, input hash, and export hash rather than by a fictitious shared seed. Each primary galaxy was queried, under leave-one-repeat-out evaluation, against reference centroids built from the remaining repeats. Table 3 reports the result: the registered primary metric, the full standardized ρ-curve RMS distance, reaches strict top-1 and top-5 source retrieval of 1.000, with a median ratio of a galaxy's own-source distance to its nearest-wrong-source distance of 0.1179; this pattern is consistent across two further retrieval metrics (the combined three-channel RMS distance and the ρ-PCA-3 representation), with per-metric top-1 accuracy between 0.9769 and 1.000.

3.2 Neighbourhood stability

Beyond source identity, repeated chamber execution preserves *local atlas structure*: the median top-5 neighbour Jaccard overlap across repeats is 1.000, and the median intraclass correlation across the first three ρ principal components is 0.99985 (minimum across galaxies: 0.9997). Within-galaxy variation under unchanged repetition is therefore far below between-galaxy separation, both at the level of individual source identity and at the level of the local neighbourhood each galaxy occupies in the atlas.

Table 3: Gate B reproducibility results across five unchanged-ladder STRUC-I executions.

Representation	Top-1	Top-5	Median own/wrong ratio
Primary: full ρ-curve RMS	1.000	1.000	0.1179
Combined three-channel RMS	1.000	1.000	0.0531
ρ-PCA-3	0.9769	1.000	0.0859
Median ICC across ρ PC1–PC3: 0.99985		Median top-5 neighbour Jaccard: 1.000	

3.3 Structural identity criterion

We adopt an operational definition: a galaxy has *structural identity under fixed observation and preparation* when repeated unchanged preparations return the same response morphology much more closely than the response morphology of any other galaxy, as quantified by the own-source-to-nearest-wrong-source distance ratio. Table 3 shows this criterion is met with wide margin for the primary sample; this criterion, and not mere numerical reproducibility of a single scalar, is what we mean by structural identity throughout the remainder of the paper.

4 Channel-Specific Architecture of the Atlas

4.1 Dimensionality of ρ, ν, and A_κ

The three channels are not interchangeable, and must not be treated as redundant projections of one another. Table 4 reports the variance retained in the first three principal components, computed separately per channel and for the combined three-channel signature. $\rho(\kappa)$ and $\nu(V_\kappa)(\kappa)$ are strongly compressible (90.8% and 88.5%), while the A_κ deficit is substantially more heterogeneous (55.9%) and does not reduce to a simple low-dimensional curve family. The combined signature retains 61.0% in three components, intermediate between the compact and heterogeneous channels, as expected from a concatenation of channels with different intrinsic dimensionality.

4.2 Atlas geometry

Figure 2b shows the combined structural atlas in its first two principal components. The map contains a dense principal population together with pronounced structural outliers, rather than one homogeneous continuum; the Weak Persistence galaxies (Section 5) occupy a visually and numerically distinct region at large positive structural-PC1.

4.3 Channel asymmetry

The three channels should be read as a *division of structural labour* rather than three noisy estimates of one quantity: ρ and ν each measure compact, low-dimensional aspects of the response, while A_κ measures a richer, less compressible aspect. This asymmetry recurs throughout the paper, from the atlas geometry reported here to the external transfer channel decomposition of Section 10.4, where the three channels again behave unequally rather than interchangeably.

4.4 Structural outliers

Galaxies occupying extreme atlas positions are disproportionately, though not exclusively, members of the Weak Persistence regime (Section 5); representative full fingerprints for a central Stable

Table 4: PCA variance retained (first three components) and structural-regime statistics.

Channel	PC1	PC2	PC3	Cumulative (3 PCs)
$\rho(\kappa)$	64.1%	17.6%	9.0%	90.8%
$\nu(V_\kappa)(\kappa)$	57.0%	24.0%	7.5%	88.5%
A_κ deficit	23.9%	17.5%	14.4%	55.9%
Combined (3 channels)	23.7%	20.2%	17.1%	61.0%
Median ρ-PC1 / structural-PC1, Stable Structure (110/121): -2.109 / -0.794				
Median ρ-PC1 / structural-PC1, Weak Persistence (11/121): $+11.218$ / $+3.076$				

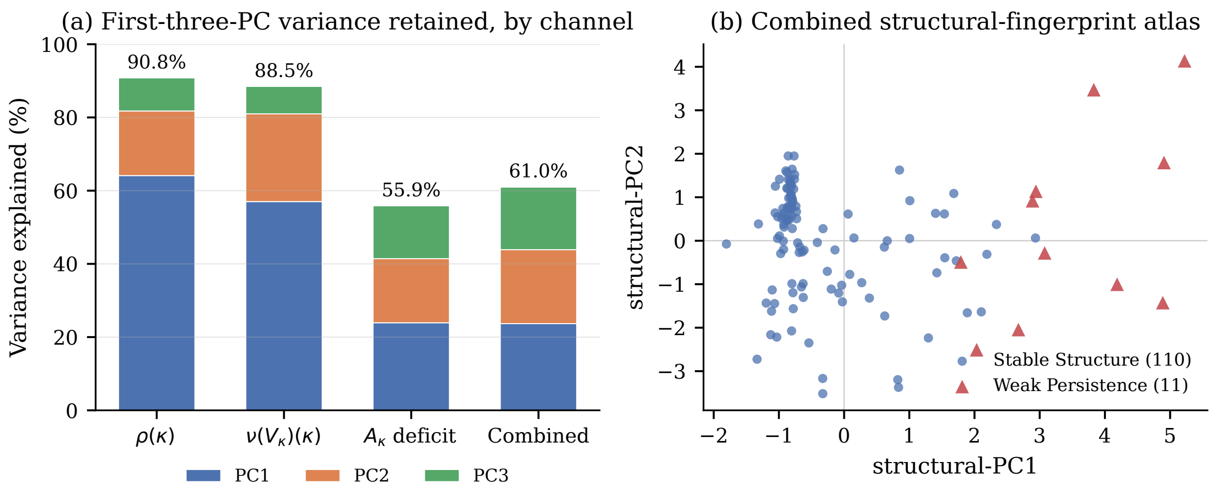

Figure 2: (a) Variance retained in the first three principal components, by channel. ρ and ν compress into 2–3 coordinates; the A_κ deficit is more heterogeneous. (b) Combined structural-fingerprint atlas (structural-PC1 vs. structural-PC2) for all 121 galaxies, coloured by persistence state.

Structure galaxy, a Stable Structure outlier, and a Weak Persistence galaxy are given in Fig. 11 (Appendix A). Stable Structure outliers show that continuous within-regime variation and discrete regime separation are simultaneously present in the atlas, a point developed further in Section 5.

5 Structural Regimes

5.1 Stable Structure and Weak Persistence

The primary sample divides into 110 Stable Structure and 11 Weak Persistence galaxies. Weak Persistence is not merely a residual label; it occupies a distinct response-space region: median ρ-PC1 is $+11.22$ versus -2.11 for Stable Structure galaxies, and median combined structural-PC1 is $+3.08$ versus -0.79 (Table 4, Fig. 3).

The Weak Persistence regime contains the 11 galaxies listed in Table 5. Their explicit identification is provided to make the regime directly usable for subsequent comparative, observational-normalization, and regime-conditioned analyses.

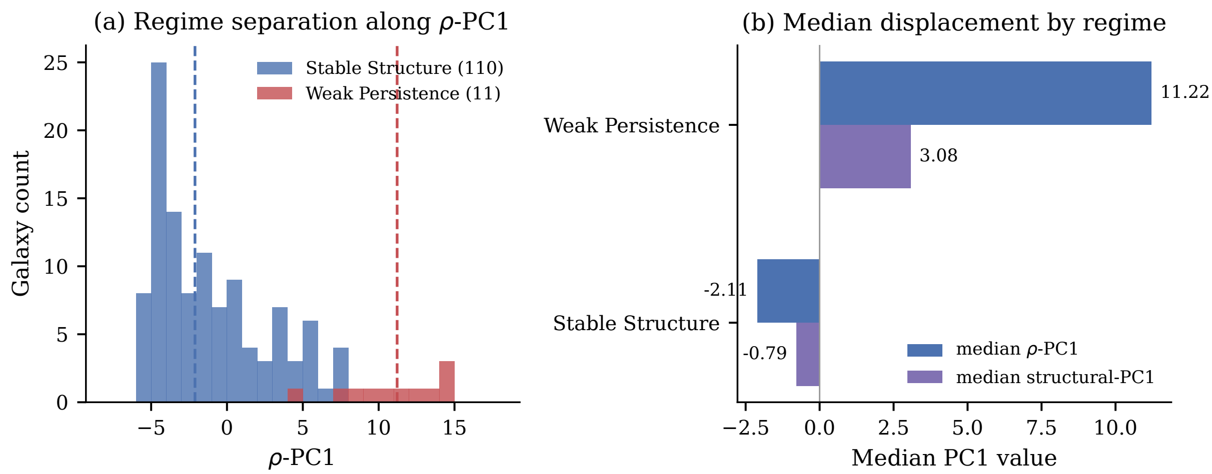

Figure 3: (a) Distribution of ρ-PC1 by persistence state, with dashed lines marking group medians. (b) Median ρ-PC1 and structural-PC1 displacement between Stable Structure and Weak Persistence galaxies.

Table 5: Galaxies classified in the validated Weak Persistence regime.

Galaxy	Galaxy
ESO563-G021	NGC 3972
F574-1	NGC 5985
IC 2574	NGC 7793
NGC 1003	UGC 02487
NGC 2998	UGC 11820
NGC 3198	

All remaining 110 galaxies in the primary sample are classified as Stable Structure under the registered STRUC-I protocol.

5.2 Regime separation in PCA space

The dominant atlas axis is therefore doing at least two things simultaneously: distinguishing individual curve shapes within a regime, and separating a discrete Weak Persistence regime from the main Stable Structure population. Both continuous galaxy-to-galaxy differentiation and discrete chamber-state organization are present in the same low-dimensional coordinate.

5.3 Identity within and between regimes

Because regime membership contributes materially to atlas position, apparent galaxy "identity" distances in the full sample can partly reflect distances between chamber regimes rather than distances between individual galaxy shapes. We term this *regime-conditioned identity*: a stability or neighbour audit should evaluate identity both globally and within the same structural regime, so that regime separation and within-regime galaxy differentiation are not conflated. All Gate B reproducibility results in Section 3 are unaffected by this distinction, since retrieval is evaluated per galaxy against the full reference gallery; regime-conditioning becomes relevant for neighbourhood and physical-alignment analyses conducted below.

5.4 Representative fingerprints

Figure 11 (Appendix A) shows complete $\rho(\kappa)$, $\nu(V_\kappa)(\kappa)$, and $A_\kappa(\kappa)$ curves for three representative galaxies: a central Stable Structure galaxy (NGC 4217), a Stable Structure outlier (UGC 11455), and a Weak Persistence galaxy (UGC 11820, whose ρ-PC1 closely matches the Weak Persistence regime median). The Weak Persistence curve shows earlier, larger-amplitude ρ growth and a deeper, broader A_κ deficit than either Stable Structure curve, consistent with its atlas position.

6 Admissibility Clusters Within the Structural Atlas

The regime analysis of Section 5 already shows that the atlas is not one undifferentiated cloud: it contains a dominant Stable Structure basin and a sharply displaced Weak Persistence regime. We now test, with explicit unsupervised clustering rather than visual inspection or the registered regime labels alone, whether the atlas supports reproducible structure beyond this two-way division: internal admissibility families within the Stable Structure basin, and their alignment with physical properties and with the persistence–fragmentation–connectivity triad of Section 7.4.

6.1 Cluster-space definition and methods

We define two complementary cluster-space representations built exclusively from structural-response information; conventional galaxy variables are retained strictly as external validation variables (Section 6.5), never as clustering inputs. **Representation A** is the full standardized multichannel curve: for each galaxy, the 40-point $\rho_g(\kappa)$, $D_{A,g}(\kappa)$, and $\tilde{\nu}_g(\kappa)$ curves (Eq. (3)) are concatenated into a 120-dimensional vector, and each of the 120 coordinates is standardized to zero mean and unit variance across the sample. **Representation B** is the combined structural-PCA coordinate (structural-PC1–PC3, Section 4.1), similarly standardized. Representation B is the primary clustering space used throughout this section, consistent with its use elsewhere in the atlas (Figs. 2 and 3); Representation A is used as an independent robustness check.

Following the recommendation that no single clustering algorithm be trusted in isolation, we apply three complementary methods, but not identically to both representations. Ward-linkage agglomerative hierarchical clustering (Euclidean distance) is applied to both Representation A and

Representation B, as an explicit representation-robustness check (Sections 6.2 and 6.3). Gaussian-mixture modelling with full covariance and Bayesian-information-criterion (BIC) model selection, and HDBSCAN, are restricted to Representation B: a full-covariance mixture fit directly in the 120-dimensional Representation A would require far more covariance parameters per component than the sample provides (110–121 galaxies), making its BIC and partition unreliable at that dimensionality, so we do not report full-covariance Gaussian-mixture or HDBSCAN results for Representation A anywhere in this section. We report silhouette, Calinski–Harabasz, and Davies–Bouldin scores for hierarchical and mixture solutions, BIC curves for mixture model selection, and adjusted Rand index (ARI) and normalized mutual information (NMI) for agreement between methods, representations, and the registered regime labels.

6.2 Global cluster structure

Applying Ward clustering to the full 121-galaxy atlas in Representation B at $k = 2$ recovers a two-cluster partition with moderate-to-strong, but incomplete, concordance with the registered Stable Structure/Weak Persistence classification (ARI = 0.678, NMI = 0.483). The smaller, displaced cluster (13 galaxies) contains 9 of the 11 Weak Persistence galaxies together with 4 Stable Structure outliers (NGC 2403, NGC 3877, UGC 06787, UGC 07524); 2 Weak Persistence galaxies (NGC 2998, UGC 02487) instead fall inside the main, 108-galaxy cluster (Fig. 4). Gaussian-mixture clustering at $k = 2$ (ARI = 0.115) and HDBSCAN (ARI = 0.000; 4 dense clusters plus 53 unassigned points) recover the regime split far more weakly, and global BIC-optimal Gaussian-mixture model selection favors $k = 5$ rather than $k = 2$, indicating that the atlas supports richer global structure than a single binary partition.

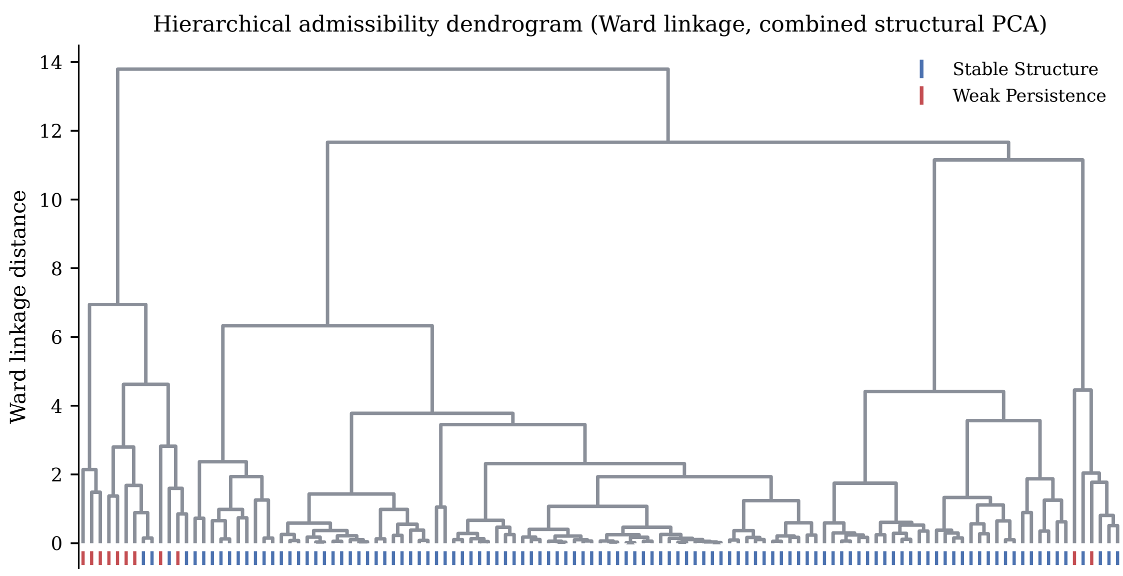

Figure 4: Ward-linkage hierarchical dendrogram for all 121 galaxies in the combined structural-PCA representation (Representation B), with leaf ticks coloured by registered persistence state. Nine of the eleven Weak Persistence galaxies form a distinct high-level branch; two (NGC 2998, UGC 02487) do not.

Clustering on the full standardized multichannel curves (Representation A) gives a materially different global $k = 2$ split: its agreement with the Representation B split is statistically negligible

(ARI ≈ -0.004). Representation A's dominant axis instead divides the Stable Structure population roughly in half (53 versus 57 galaxies) without isolating Weak Persistence as a class – all 11 Weak Persistence galaxies fall inside the larger, 68-galaxy half. This indicates that the raw, unreduced curves are shaped by sources of variance other than persistence regime, plausibly related to the same observational sampling-geometry sensitivity already identified through Gate C (Section 9.3), rather than by a defect in either representation. We conclude that the visually separated 110/11 structure of Section 5 is recovered by unsupervised clustering only partially and representation-dependently: a real, low-dimensional displacement along the dominant combined-PCA axis, but not a universally density-separated cluster across every representation and algorithm tested.

6.3 Within-regime admissibility clusters

Because the large Weak Persistence displacement can dominate a global partition (Section 6.2), we repeat the analysis on the 110 Stable Structure galaxies alone, in Representation B. Ward-linkage silhouette scores increase only slowly, without a sharp elbow, across $k = 2$–6 (0.366, 0.395, 0.393, 0.402, 0.408); Gaussian-mixture BIC is minimized at $k = 3$ (BIC = 380.7, versus 464.3 at $k = 2$ and 407.5 at $k = 4$); HDBSCAN identifies 4 dense sub-clusters together with 42 unassigned galaxies (38% of the Stable Structure population). No single algorithm identifies one unambiguous cluster count. Representation A gives a more modest within-regime signal (best Ward silhouette 0.286, at $k = 2$) but agrees well with Representation B at matched $k = 2$ (ARI = 0.928) – unlike the global comparison of Section 6.2, the two representations describe compatible structure once the Weak Persistence displacement is removed.

6.4 Cluster stability and consensus construction

To convert this multi-method, multi-k evidence into one defensible partition, we build a consensus solution. An ensemble of seven base clusterings (Ward at $k = 3, 4, 5$; Gaussian-mixture at $k = 3, 4, 5$; one HDBSCAN run) is combined with 200 subsample replicates, each an independent random 85% subsample of the 110 Stable Structure galaxies refit with both Ward and Gaussian-mixture clustering at a reference $k = 4$, contributing a further 400 clustering solutions. The full ensemble is used to build a pairwise co-assignment matrix: for each pair of galaxies, the fraction of clustering solutions containing both galaxies in which they are placed in the same cluster (Fig. 5). Co-assignment is calculated only across solutions in which both galaxies are present and assigned a cluster label: a subsample replicate contributes to the denominator only for the galaxy pairs it contains, and HDBSCAN noise points are treated as unassigned for that run rather than as members of a shared noise cluster. Average-linkage clustering of the resulting co-assignment distance (1−co-assignment fraction) is cut at the k that maximizes silhouette width on that distance; this is maximized at $k = 5$ (silhouette 0.825, versus 0.778–0.789 at $k = 3$, 4, 6), which we adopt as the consensus admissibility-cluster count reported in Table 6. Clusters are numbered 0–4 in order of increasing median ρ-PC1 for readability; the numeric labels carry no other meaning and are not guaranteed to reproduce under a different consensus run.

Bootstrap subsample stability, measured as the adjusted Rand index between each 85%-subsample clustering and the reference $k = 4$ clustering restricted to the sampled galaxies (200 replicates), is moderate-to-good: mean ARI = 0.694 ± 0.120 for Ward and 0.742 ± 0.125 for Gaussian mixtures. Cross-method agreement at the consensus $k = 5$ is more variable: Ward and Gaussian-mixture partitions agree only moderately (ARI = 0.536), while Ward and HDBSCAN agree strongly on the galaxies HDBSCAN does assign (ARI = 0.881). Sensitivity to PCA dimension is also moderate: Ward clustering on structural-PC1–2 versus PC1–3 gives ARI = 0.704, and PC1–3 versus the 9-

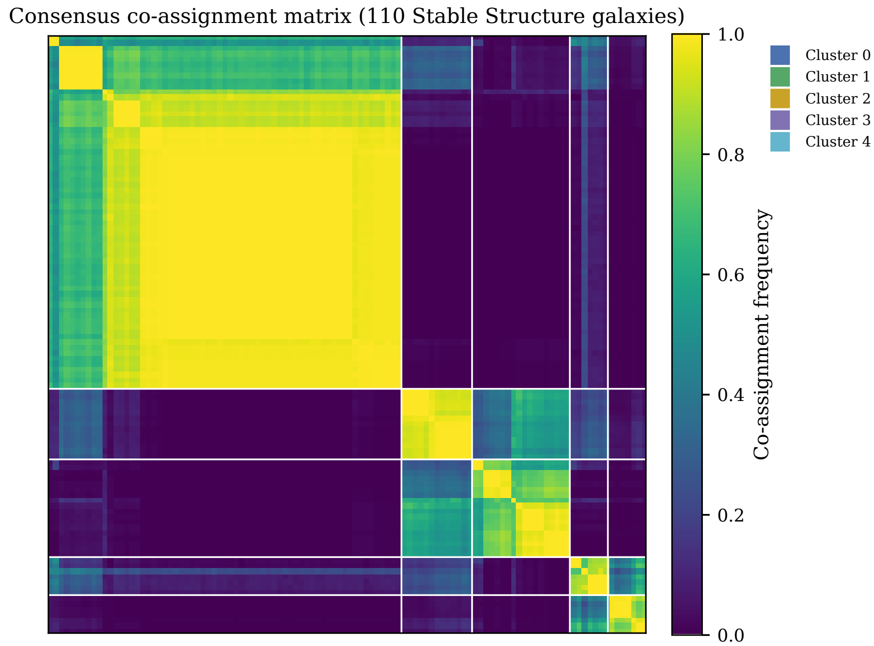

Figure 5: Consensus co-assignment matrix for the 110 Stable Structure galaxies, ordered by average-linkage clustering of the co-assignment distance and annotated by final consensus cluster. Bright, block-diagonal structure indicates strong internal reproducibility of each admissibility cluster across the full ensemble of methods, k values, and resampling replicates.

dimensional per-channel PCA space (Section 4.1) gives ARI = 0.648. Each final consensus cluster nonetheless shows strong internal coherence in the co-assignment matrix itself: within-cluster mean co-assignment ranges from 0.819 to 0.967, versus 0.037–0.144 between clusters (Table 6). Because Gate B already demonstrates near-perfect unchanged-preparation reproducibility of the underlying chamber curves (top-1/top-5 retrieval 1.000, median ICC 0.99985; Section 3), the residual clustering instability reported here is attributable to the clustering procedure's sensitivity to algorithm, subsampling, and representation, not to non-reproducibility of the underlying per-galaxy chamber signatures.

We regard the $k = 5$ consensus clusters as reproducible with respect to resampling and largely, though not perfectly, reproducible across clustering algorithm and PCA dimension. Beyond the registered Stable Structure/Weak Persistence division, whose displacement along the dominant combined-PCA axis is clear but whose unsupervised recovery is representation- and method-dependent rather than robustly and near-unanimously recovered (Section 6.2), the precise number and boundaries of finer admissibility clusters remain method-sensitive. Table 6 is therefore best read as the atlas's most defensible current candidate taxonomy of reproducible admissibility families within the Stable Structure basin, not as a final or uniquely correct partition.

6.5 Physical and cross-chamber validation

Kruskal–Wallis tests across the five consensus clusters show significant differences in Hubble type ($H = 11.30$, $p = 0.023$), flat rotation velocity ($H = 11.34$, $p = 0.023$), 3.6 μm luminosity ($H = 15.47$, $p = 0.0038$), effective surface brightness ($H = 17.36$, $p = 0.0016$), H I mass ($H = 10.98$, $p = 0.027$), and median ladder gap ($H = 17.90$, $p = 0.0013$), but not in valid radial-point count ($H =$

Table 6: Consensus admissibility-cluster registry within the Stable Structure regime ($k = 5$, $n = 110$). Full Percolation fraction is the STRUC-PERC-I connectivity phenotype C_g within each cluster; the sample-wide Stable Structure rate is 15/110 (13.6%).

Cluster	Size	Median ρ-PC1	Full Percolation	Physical character
0	65	−4.07	3/65 (4.6%)	Bright, high-V_{flat}, high-surface-brightness core
1	13	+0.67	3/13 (23.1%)	Late-type, low-luminosity, low-surface-brightness family
2	18	+0.96	8/18 (44.4%)	Intermediate physical scale; markedly enriched in Full Percolation
3	7	+3.31	0/7 (0%)	Small family; uniformly Hard Fragmentation
4	7	+5.59	1/7 (14.3%)	Most ρ-PC1-displaced Stable Structure family, approaching the Weak Persistence direction

Full galaxy membership by cluster is given in Table 13 (Appendix A).

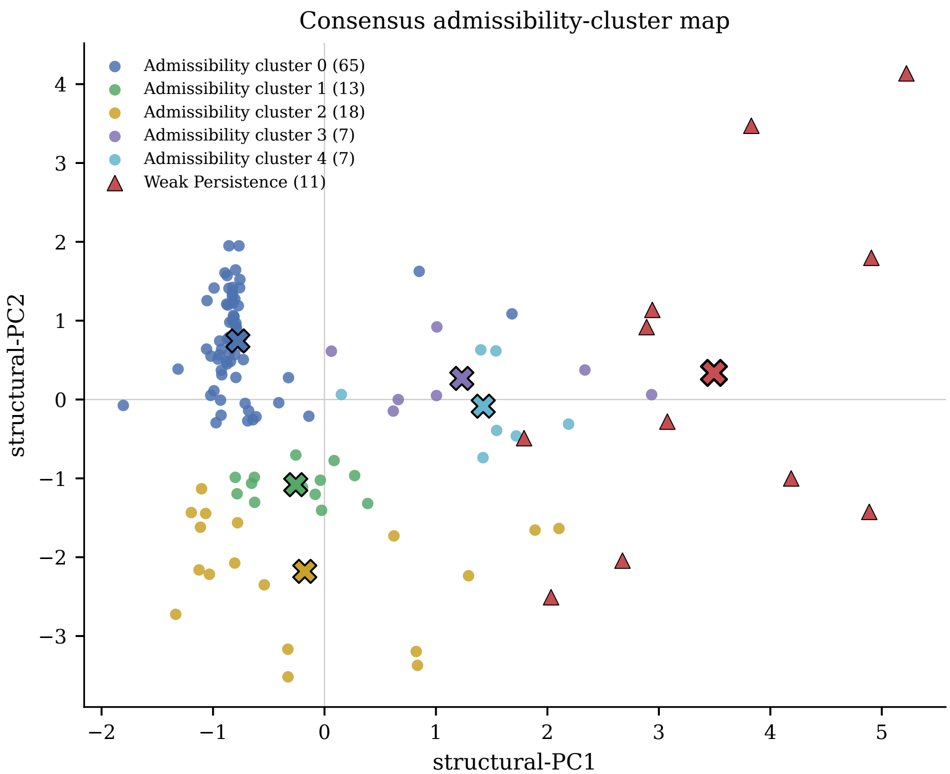

Figure 6: Consensus admissibility-cluster map: combined structural atlas (structural-PC1 vs. structural-PC2) for all 121 galaxies, coloured by consensus admissibility cluster within Stable Structure (circles) or marked as Weak Persistence (red triangles); large X markers show cluster centroids.

7.48, $p = 0.113$). Applying a Benjamini–Hochberg false-discovery-rate correction across these seven tests at $q = 0.05$ leaves all six physical and ladder-gap differences significant (adjusted $p \leq 0.031$) and valid radial-point count non-significant (adjusted $p = 0.113$), so the physical-alignment result is not an artifact of testing seven properties at once. The clusters are therefore substantially, but not purely, physically organized: their separation is not simply an artifact of differing observed point counts, though it is not fully independent of observational sampling geometry either (Section 9.3). Figure 6 and Table 6 summarize the resulting physical character: Cluster 0, the dominant admissible family, contains the brightest, highest-V_{flat}, highest-surface-brightness galaxies (medians: Hubble type 5, $V_{\mathrm{flat}} = 154.0\ \mathrm{km\,s^{-1}}$, luminosity $68.6 \times 10^9\,L_\odot$, surface brightness $421.5\,L_\odot\,\mathrm{pc}^{-2}$); Cluster 1 is the latest-type, faintest family (median Hubble type 9, luminosity $2.9 \times 10^9\,L_\odot$); Clusters 3 and 4 are the smallest and most structurally displaced sub-families, with Cluster 4 reaching the highest median ρ-PC1 among all Stable Structure clusters (+5.59), approaching, without crossing into, the Weak Persistence direction of the atlas (median ρ-PC1 = +11.22; Section 5.1).

Cross-tabulating consensus cluster membership against the STRUC-PERC-I connectivity verdict – the connectivity component C_g of the persistence–fragmentation–connectivity triad (Section 7.4) – shows that Full Percolation is not evenly distributed across the admissible families (Table 6): Cluster 2 is markedly enriched (8/18 = 44.4% Full Percolation, versus 17/121 = 14.0% over the full primary sample and 15/110 = 13.6% within Stable Structure overall), while Cluster 3 shows no Full Percolation (0/7) and the remaining clusters lie closer to the sample-wide rate. An omnibus test of consensus-cluster membership against Full Percolation status (the 5×2 contingency table underlying Table 6) is significant ($\chi^2 = 21.09$, df = 4, $p = 3.0 \times 10^{-4}$), though several clusters have small expected cell counts under that approximation; the more targeted and exact test is a one-sided Fisher's exact test of Cluster 2 (8/18 Full Percolation) against the pooled remaining Stable Structure clusters (7/92 Full Percolation), which gives $p = 3.6 \times 10^{-4}$ (odds ratio 9.7). Admissibility clustering, built solely from STRUC-I response geometry, therefore recovers a subgroup that is independently, and statistically strongly, distinguished by its STRUC-PERC-I connectivity behaviour – a cross-chamber phenotype alignment consistent with, and extending, the dissociability of persistence and connectivity already established in Section 7.2.

6.6 Atlas hierarchy: regimes, clusters, and structural analogues

These results motivate a three-level hierarchy within the atlas, extending the regime/neighbour language used elsewhere in this paper. A *structural regime* (Stable Structure or Weak Persistence, Section 5) is a broad response state. An *admissibility cluster* (this section) is a reproducible family of galaxies occupying a shared region of multichannel response space within a regime. A *structural analogue* (Section 8.3) is a close galaxy pair that may belong to the same admissibility cluster despite differing conventional classification. Symbolically,

$$\text{Atlas} \supset \text{Regimes} \supset \text{Admissibility clusters} \supset \text{Structural neighbours.} \tag{5}$$

The two members of every listed cross-class structural-analogue pair (Table 9) fall within the same consensus admissibility cluster: DDO 168/NGC 3893, NGC 5585/NGC 7814, F563-1/UGC 06614, and UGC 00128/UGC 06614 are each internal to Cluster 0, and NGC 0024/DDO 064 are both internal to Cluster 2. This is independent confirmation that these pairs reflect genuine shared admissibility-cluster membership, discovered without reference to their pairwise distances, rather than isolated nearest-neighbour coincidences.

Table 7: Cross-chamber persistence, fragmentation, and connectivity counts ($n = 121$).

Structural state	Count
Stable Structure	110
Weak Persistence	11
Hard Fragmentation	103
Stable Structure *and* Hard Fragmentation	94
Full Percolation	17
Full Percolation reached only via adaptive extension (κ-connect > 1)	17 / 17

7 Persistence, Fragmentation, and Connectivity

7.1 Cross-chamber result

Combining STRUC-I persistence with STRUC-PERC-I connectivity for the same 121 galaxies gives a result that is, on its own, one of the paper's strongest conceptual findings (Table 7, Fig. 7): 110/121 galaxies are classified Stable Structure; 103/121 show Hard Fragmentation; 94/121 are *simultaneously* Stable Structure and Hard Fragmentation; and only 17/121 reach Full Percolation, in every case only through the adaptive κ-connect extension, with none reaching full connectivity on the native κ grid.

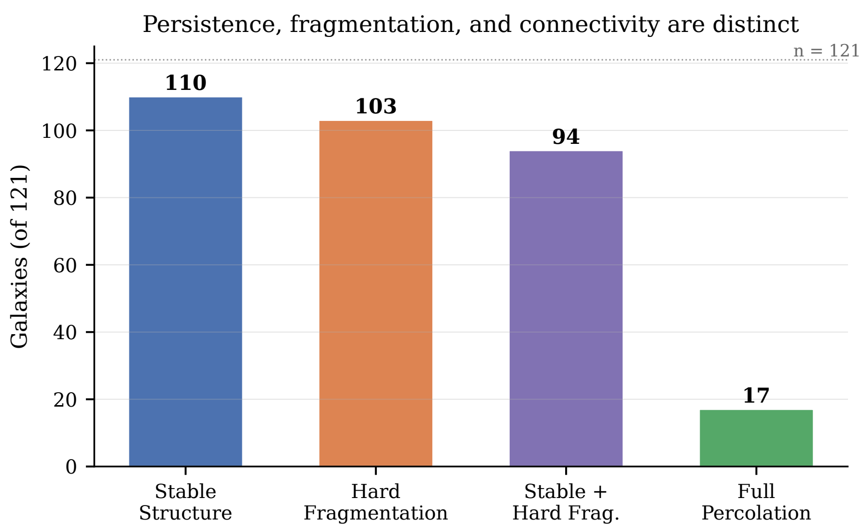

Figure 7: Persistence, fragmentation, and connectivity counts across the 121-galaxy primary sample. Structural persistence and hard fragmentation frequently coexist; full percolation is comparatively rare and, when reached, requires the adaptive κ extension.

7.2 Distinct and dissociable structural dimensions

These counts demonstrate that persistence is not connectivity, and fragmentation is not structural failure. A galaxy can retain a stable multiscale response, remain fragmented, fail to percolate on the native κ interval, and reach full connectivity only through an adaptive extension. We find no detectable association between Stable/Weak Persistence and whether full percolation is eventually reached, reinforcing that these are empirically distinct and dissociable structural observables rather than two views of a single stability property.

7.3 Adaptive connectivity

We interpret κ-connect > 1 as an extension requirement of the connectivity measurement, not as a failure of structure: the same ladder that requires an adaptive extension to percolate can simultaneously show a highly stable admissibility response under STRUC-I. Connectivity and persistence are answering different structural questions about the same ordered ladder.

7.4 Three-coordinate structural description

We formalize this as a three-axis structural description,

$$\mathcal{S}_g = \big(P_g,\, \Phi_g,\, C_g\big), \tag{6}$$

where P_g is structural persistence, Φ_g is fragmentation behaviour, and C_g is connectivity or percolation reach. We use Φ_g rather than F_g to avoid collision with the fingerprint symbol $\mathcal{F}_g(\kappa)$. This triad is more informative than a single stable/unstable verdict, and Table 7 shows that its three components are empirically dissociable within the primary sample. Connectivity cannot be used as a synonym for structural stability or admissibility.

8 Physical Organization and Structural Analogues

8.1 Correlations with galaxy properties

Some structural coordinates organize galaxies according to ordinary physical scale. Figure 8a reports Spearman correlations between structural-PC2 and four conventional properties: effective surface brightness ($\rho_s = 0.463$), 3.6 μm luminosity ($\rho_s = 0.456$), flat rotation velocity ($\rho_s = 0.404$), and H I mass ($\rho_s = 0.391$; all $p < 10^{-4}$, $n = 121$). At least one major atlas direction is therefore organizing galaxies by broad physical scale, mass, and surface-brightness characteristics, in addition to whatever purely structural information it encodes.

8.2 Nearest-neighbour physical coherence

Nearest structural neighbours are more physically alike than typical galaxy pairs, though matches are far from exact. Table 8 compares median property differences for structural nearest-neighbour pairs against the all-galaxy-pair median: Hubble type differs by a median of 2 for structural neighbours versus 3 for all pairs; flat velocity by 62.7 versus 80.1 $\mathrm{km\,s^{-1}}$; 3.6 μm luminosity by 37.3 versus $65.6 \times 10^9\, L_\odot$; effective surface brightness by 275.1 versus $396.1\, L_\odot\,\mathrm{pc}^{-2}$; and H I mass by 2.41 versus $3.80 \times 10^9\, M_\odot$. Figure 8b expresses these as a percentage reduction relative to the all-pair baseline, ranging from 22% (flat velocity) to 43% (3.6 μm luminosity).

8.3 Cross-class analogues

Because structural-neighbour matches are systematically closer than the all-galaxy-pair baseline, while remaining far from exact, the atlas also identifies *structural analogues*: galaxy pairs that are close in structural-response space despite differing conventional classifications. This is not a byproduct to be explained away. It is a second, independent form of scientific content in the atlas,

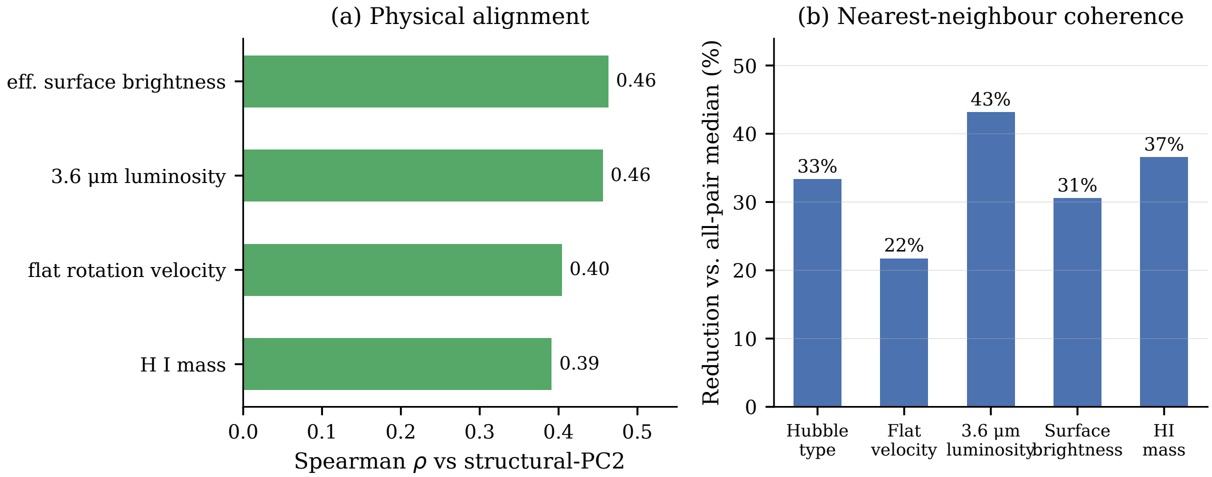

Figure 8: (a) Spearman correlation between structural-PC2 and four conventional physical properties. (b) Percentage reduction in median property difference for structural nearest-neighbour pairs relative to the all-galaxy-pair median, across five conventional properties.

Table 8: Structural-neighbour coherence versus the all-galaxy-pair baseline.

Property	Structural neighbour	All galaxy pairs	Reduction
Hubble type	2	3	33%
Flat velocity ($\mathrm{km\,s^{-1}}$)	62.7	80.1	22%
3.6 μm luminosity ($10^9\,L_\odot$)	37.3	65.6	43%
Effective surface brightness ($L_\odot\,\mathrm{pc}^{-2}$)	275.1	396.1	31%
H I mass ($10^9\,M_\odot$)	2.41	3.80	37%

complementary to its preservation of conventional physical organization. Table 9 names five such pairs, each a nearest structural neighbour despite a Hubble-type difference of at least 3 and a flat-velocity difference of tens to over 100 $\mathrm{km\,s^{-1}}$.

8.4 What the atlas adds

The atlas does not replace conventional galaxy classification; Sections 8.1 and 8.2 show that it substantially preserves it. What the atlas supplies in addition is a new relation between galaxies: similarity of complete multiscale structural response, which is neither reducible to standard morphological class nor to any single conventional physical parameter.

Table 9: Cross-class structural analogues: nearest structural-neighbour pairs with differing conventional classification.

Galaxy A	Galaxy B	Struct. dist.	Hubble types	$\Delta V_{\rm flat}$ ($\rm km\,s^{-1}$)	Why notable
DDO 168	NGC 3893	0.079	10 vs. 5	120.6	Dwarf irregular and massive spiral, nearly identical structural response.
NGC 5585	NGC 7814	0.080	7 vs. 2	128.6	Late-type disk and early-type bulge-dominated spiral, close structural neighbours.
NGC 0024	DDO 064	0.135	5 vs. 10	60.2	Spiral and dwarf irregular pairing, opposite ends of the Hubble sequence.
F563-1	UGC 06614	0.144	9 vs. 1	89.9	Low-surface-brightness dwarf paired with an early-type high-velocity disk; UGC 06614 also shows the largest combined transfer benefit in Section 10.5.
UGC 00128	UGC 06614	0.159	8 vs. 1	70.5	A second, independent late-type pairing with the same high-velocity structural analogue.

9 The Observational Transformation Layer

9.1 Gate C design

A structural fingerprint is observed through a measurement channel, and changes in that channel transform its recorded coordinates. Gate C probes this transformation directly through four branches applied to the same primary sample: leave-one-radius-out deletion; deterministic n=10 thinning; fixed-count radial-coverage modification; and disk/bulge mass-to-light sensitivity.

9.2 Transformation strength

Table 10 and Fig. 12 (Appendix A) report top-1 and top-5 retrieval and the median own-to-nearest-wrong ratio for each branch. None of the four branches retains the Gate B unchanged-repetition thresholds: retrieval accuracy falls from 1.000 under unchanged repetition to between 0.050 and 0.551 (top-1) depending on branch, and the median own/nearest-wrong ratio rises correspondingly.

9.3 Hierarchy of sensitivity

The four branches should be read as different strengths and kinds of transformation, not as four uniform failures. Leave-one-radius-out and mass-to-light sensitivity leave retrieval performance comparatively high (top-1 $\geq$ 0.52); deterministic thinning and, especially, fixed-count coverage change produce substantially larger transformation (top-1 as low as 0.050). This ordering is itself informative: it indicates that radial coverage and point count are more active coordinates in fingerprint expression than single-point deletion or reasonable mass-to-light variation. Consistent with this, Section 8 independently found that final ρ correlates with median ladder gap ($\rho_s = -0.541$), structural-PC3 correlates with valid radial-point count ($\rho_s = -0.493$), and vulnerability fraction

Table 10: Gate C representation-response results by branch.

Branch	Top-1	Top-5	Median own/wrong ratio	Geometry-matched top-1
Leave one radius out	0.551	0.730	0.916	0.726
Deterministic n=10 thinning	0.165	0.314	1.835	0.347
Fixed-count coverage	0.050	0.198	2.138	0.256
Mass-to-light sensitivity	0.517	0.619	0.916	0.649

Geometry-matched top-1 restricts the candidate reference pool to each galaxy's ten nearest neighbours in ordinary sampling-geometry space (ladder size, gap statistics, and radial-coverage descriptors, independent of the structural-response channels) rather than the full 121-galaxy gallery. It tests whether source retrieval survives when compared only against galaxies with similar observational geometry, rather than succeeding merely because of gross differences in sampling density or coverage.

correlates with median ladder gap ($\rho_s = +0.455$; all $p < 10^{-4}$). The atlas therefore carries an observational transformation layer in addition to its physical layer, formally

$$\mathcal{F}_g^{\rm obs} = \mathcal{T}_{\rm obs}\left[\mathcal{F}_g^{\rm gal}\right], \tag{7}$$

where $\mathcal{F}_g^{\rm gal}$ is the galaxy-associated structural response and $\mathcal{T}_{\rm obs}$ is the observational and preparation transformation. Deletion, thinning, coverage change, and baryonic preparation do not make the fingerprint unreal; they measure the strength and character of $\mathcal{T}_{\rm obs}$.

9.4 Toward normalized fingerprints

The present programme maps the existence and strength of $\mathcal{T}_{\rm obs}$; it does not yet invert it. The natural next objective is a normalized fingerprint

$$\hat{\mathcal{F}}_g = \mathcal{N}\left(\mathcal{F}_g^{\rm obs},\, \mathcal{O}_g\right), \tag{8}$$

where $\mathcal{O}_g$ records observational geometry (radial coverage, point count, ladder spacing) and $\mathcal{N}$ is a normalization mapping designed to remove the observational layer while preserving the galaxy-associated response. Gate C's four branches, and the geometry correlations of Section 8, together define the empirical response surface that such a normalization must fit.

10 External Transfer Test

10.1 Why test transfer, and against what baseline

A real structural fingerprint should be tested beyond self-identification and atlas organization. Section 1.2 introduced the RAR as the principal empirical baseline for local baryon–dynamics coupling; we now use it, together with progressively richer local and whole-galaxy baselines, as the null hypothesis for whether chamber signatures carry *additional*, held-out predictive information for gravitational discrepancy. Define the discrepancy coordinate

$$\Xi(R) = \frac{g_{\rm obs}(R)}{g_{\rm bar}(R)} - 1, \qquad \text{or, in the regression protocol,} \qquad y(R) = \log\left(\frac{g_{\rm obs}(R)}{g_{\rm bar}(R)}\right). \tag{9}$$

Because the RAR already maps local baryonic acceleration to local dynamical discrepancy, a structural chamber signature is gravitationally informative only if it predicts held-out radial behaviour

Table 11: Static whole-galaxy model hierarchy used in the external transfer test.

Model	Experimental role
M0	Local g_{bar} only; directly RAR-like baseline.
M1	Stronger local baseline: g_{bar} plus normalized radius.
M2	Local baseline plus conventional whole-galaxy ladder statistics.
M-I	M2 plus STRUC-I compact signatures.
M-P	M2 plus STRUC-PERC-I compact signatures.
M-IP	M2 plus both compact signature families.
Full-curve models	Fold-local curve representations, decomposed by channel (ρ, A_κ, ν, all channels).

beyond the information already supplied by g_{bar} and ordinary, non-UNNS controls. This gives the model hierarchy of Table 11: M0 represents the RAR operationally through a cubic-spline function of $\log_{10} g_{\text{bar}}$, rather than imposing the particular exponential interpolation of Eq. (2); M1 strengthens this local baseline with normalized radius; M2 adds conventional whole-galaxy ladder statistics; M-I, M-P, and M-IP add STRUC-I, STRUC-PERC-I, and combined compact chamber signatures on top of M2; and the full-curve models append fold-local PCA scores derived from the complete chamber curves, with PCA fitted exclusively on training galaxies and the resulting galaxy-level scores attached to the radial rows of each corresponding galaxy. The full-curve models are therefore richer than the compact summaries, but they remain static, whole-galaxy descriptors like every other tested architecture. We emphasize that neither M1 nor M2 is a pure RAR model — M1 adds normalized radius and M2 adds whole-galaxy profile statistics — so the RAR establishes the conceptual null while M1 and M2 provide increasingly strong *operational* baselines that any chamber-augmented model must beat under galaxy-held-out cross-validation.

10.2 Static whole-galaxy model comparison

Figure 9a and Table 12 report galaxy-balanced held-out RMSE under galaxy-held-out cross-validation. The best-performing model is M1 (RMSE 0.18428), the stronger local baseline; M0 performs almost as well (0.18562). Every model that adds whole-galaxy descriptors performs worse: M2 (0.18839), M-I (0.19005), M-P (0.19168), and M-IP (0.19455, a -3.27% relative change from M2, bootstrap 95% CI $[-0.0146, 0.0022]$). This ordering shows something broader than "chamber signatures did not help": even ordinary, non-UNNS whole-galaxy profile statistics (M2) reduce cross-galaxy performance relative to the simpler local baseline. The static whole-galaxy architecture, not the chamber variables specifically, is the identified limitation.

10.3 Compact-signature overfit

The combined compact model improves its *training* fit relative to M2 (training R^2: $0.7198 \rightarrow 0.7272$) while its *held-out* fit deteriorates (held-out R^2: $0.6847 \rightarrow 0.6516$), and its calibration slope falls from 0.925 to 0.888. A better training description did not produce better held-out transfer. Feature ablation on M-IP (Fig. 9b) shows that removing the compact A_κ summaries yields the largest single improvement in held-out RMSE, from 0.19455 to 0.18981, consistent with A_κ being the most heterogeneous and least compact channel identified in Section 4.1. This behaviour is consistent with the compact chamber signatures encoding galaxy identity or sample-specific structure that a

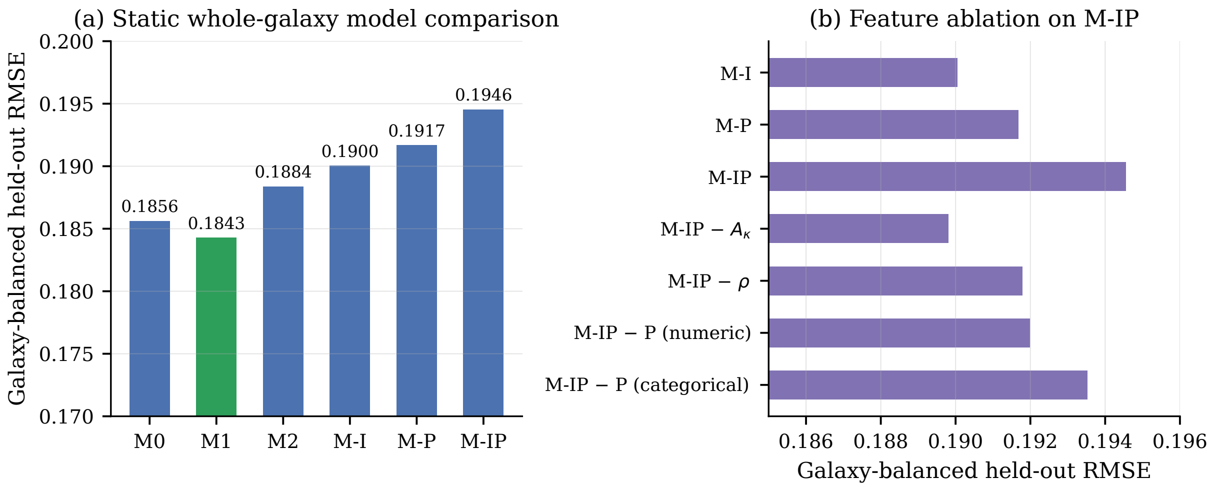

Figure 9: (a) Galaxy-balanced held-out RMSE for the static whole-galaxy model hierarchy of Table 11; the local baseline M1 (green) is the best-performing model. (b) Feature ablation on the combined M-IP model: removing A_κ improves held-out RMSE the most of any single ablation.

static, whole-galaxy model can fit but not transfer, rather than with a transferable dynamical law.

10.4 Full-curve channel decomposition

Full curves confirm identity but not a universal gravitational bridge. Correct galaxy-to-curve assignment beats shuffled assignment (one-sided $p \approx 0.020$), showing that the curves carry recoverable galaxy-specific information; nevertheless, the aggregate full-curve gain over M2 is small (+0.44%), only 1 of 5 folds improves, and the bootstrap 95% interval crosses zero. The channel decomposition (Fig. 10, Table 12) shows the channels are not interchangeable: $\rho(\kappa)$ is the leading transfer-sensitive channel, improving aggregate performance by +1.14% with positive fold-level changes in 4 of 5 folds, although it does not clear the full registered confirmation gate; $A_\kappa(\kappa)$ gives a smaller aggregate gain (+0.46%, 3/5 folds); and $\nu(V_\kappa)(\kappa)$ is actively detrimental for this target (−1.52%, bootstrap interval entirely on the detrimental side). Combining every channel is therefore counterproductive relative to using ρ alone: the all-channel result (+0.44%, 1/5 folds) is weaker than the ρ-only result because ν drags the combination down.

10.5 Galaxy-level heterogeneity

The aggregate result conceals substantial galaxy-to-galaxy heterogeneity. The compact combined model improved predictions for 53 of 121 galaxies; the full ρ curve improved predictions for 71 of 121. Several galaxies benefited consistently across representations (e.g. UGC 06614, UGC 02487, F563-V2, UGC 07125, DDO 064), while others deteriorated consistently (e.g. IC 2574, F563-1, NGC 3877, NGC 4183). Representation-dependent reversals also occur: compact signatures substantially worsen NGC 3972, while the full ρ curve improves it (Fig. 13, Appendix A). This heterogeneity is a central clue rather than noise to be averaged away: it is consistent with distinct structural regimes, profile-quality dependence, or local radial-transition phenomena that a single global linear coefficient cannot represent.

Table 12: External transfer results: static compact models and full-curve channel decomposition.

Model	Galaxy-balanced RMSE	Relative Δ vs. M2	Positive folds
M0	0.18562	—	—
M1 (best static model)	0.18428	—	—
M2	0.18839	0.00%	—
M-I	0.19005	−0.88%	—
M-P	0.19168	−1.75%	—
M-IP	0.19455	−3.27%	—
Full curves, all channels	0.18757	+0.44%	1 / 5
Full curves, $\rho(\kappa)$	0.18624	+1.14%	4 / 5
Full curves, $A_\kappa(\kappa)$	0.18753	+0.46%	3 / 5
Full curves, $\nu(V_\kappa)(\kappa)$	0.19125	−1.52%	1 / 5

10.6 Transfer boundary

The atlas is real, but the correct predictive architecture is not a single static whole-galaxy descriptor. The rejected object in this experiment is not the fingerprint; it is the specific assumption that a complete galaxy fingerprint can be inserted, as a static global descriptor, into one universal radial-prediction model. This is a positive methodological result: it closes one tested architecture while motivating local, regime-conditioned, and channel-selective successors, with $\rho(\kappa)$ identified as the leading candidate channel.

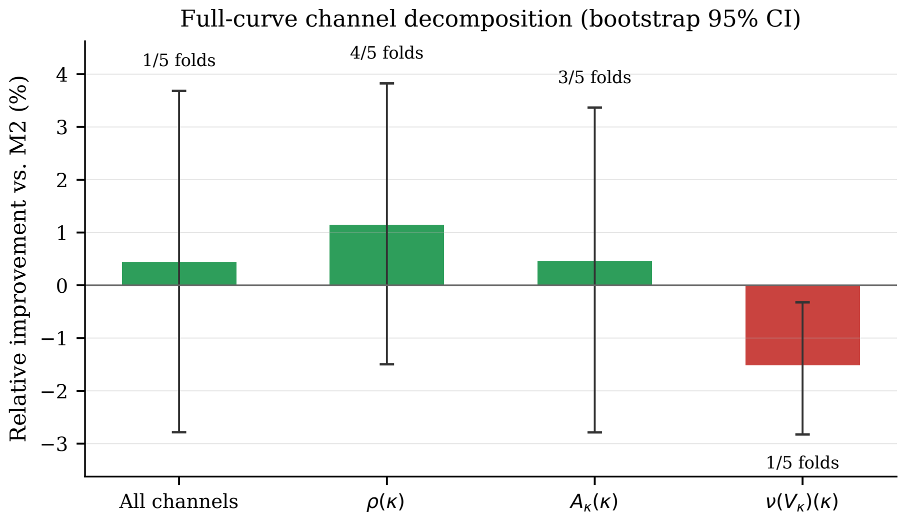

Figure 10: Full-curve channel decomposition: relative galaxy-balanced RMSE improvement over M2, with bootstrap 95% confidence intervals and the number of positive folds (of 5) annotated. $\rho(\kappa)$ is the leading transfer-sensitive channel; $\nu(V_\kappa)(\kappa)$ is detrimental.

11 Discussion

We organize the discussion around seven propositions established by the results above.

1. **Galaxy baryonic-acceleration ladders possess structural identity.** Gate B shows near-perfect unchanged-preparation source retrieval and neighbourhood stability (Section 3).

2. **Structural identity is multichannel.** ρ, ν, and A_κ carry different, complementary structural information rather than redundant estimates of one quantity (Section 4).

3. **The atlas contains both continuous differences and discrete regimes.** Stable Structure and Weak Persistence form a genuine response regime, not merely a label on eleven unusual galaxies (Section 5).

4. **Persistence, fragmentation, and connectivity are separate dimensions.** A galaxy can be structurally persistent while fragmented and non-percolating on the native κ range (Section 7).

5. **The Stable Structure basin resolves into a five-family consensus admissibility structure.** Built from Ward, Gaussian-mixture, and HDBSCAN clustering with bootstrap stability checks, the partition shows strong internal co-assignment coherence, moderate-to-good resampling stability, and statistically significant differentiation in independent physical properties and STRUC-PERC-I connectivity phenotype, extending the atlas's regime structure into a three-level regime, cluster, and neighbour hierarchy (Section 6).

6. **Observed fingerprints include an observational transformation layer.** Sampling density, radial coverage, and preparation choices are active coordinates in fingerprint expression, alongside a physical layer that preserves conventional galaxy organization (Sections 8 and 9).

7. **Predictive transfer requires a local, regime-sensitive architecture.** A single static whole-galaxy descriptor does not improve held-out radial prediction, while $\rho(\kappa)$ and galaxy-level heterogeneity motivate a local, regime-conditioned, channel-selective successor (Section 10).

11.1 Relation to conventional galaxy classification

The atlas substantially preserves conventional galaxy organization: structural neighbours differ less than the all-galaxy-pair baseline in Hubble type, luminosity, surface brightness, rotation velocity, and gas content (Section 8.2). It adds to this a relation that conventional classification does not supply: quantitative similarity of complete multiscale structural response, including structural analogues among galaxies with different conventional classifications (Section 8.3).

11.2 Structural fingerprints as a new data object

We treat the response curves, PCA coordinates, and atlas distances as reusable scientific outputs in their own right, independent of any single downstream predictive task. This framing follows directly from Section 3–Section 8: the atlas is established before, and independently of, the transfer test of Section 10.

11.3 Limits of static scalarization

Compacting a full structural trajectory into a handful of galaxy-level features, as the compact M-I/M-P/M-IP models do, can discard transition-local information that only the full curve retains. Sections 10.3 and 10.4 show this directly: the full ρ curve outperforms every compact combination in the transfer test, and it does so unevenly across galaxies (Section 10.5), consistent with locally concentrated, rather than globally averaged, structural information.

11.4 Structural identity and RAR-level predictive sufficiency are different questions

This is the most valuable use of the RAR baseline established in Sections 1.2 and 10. The chamber curves identify and organize galaxies (Section 3–Section 8); the Gravity Bridge experiment asks whether those whole-galaxy fingerprints add held-out predictive information for radial dynamical discrepancy *beyond* local baryonic acceleration (Section 10). These are different questions, and our results answer them differently without contradiction: *the RAR organizes the local baryon–dynamics relation, while the UNNS atlas organizes the multiscale structural identity of the complete canonicalized baryonic-acceleration ladder. The present static transfer experiment shows that the second organization cannot simply be appended as a global descriptor to improve the first.* This is a more precise statement than saying the chambers "failed to beat MOND," and a more accurate one than treating the fingerprint as unnecessary: the tested chamber augmentations did not demonstrate incremental held-out predictive power beyond the local and conventional baselines, but this does not establish that g_{bar} alone is a complete or uniquely sufficient explanation of galaxy dynamics — indeed M1 (g_{bar} plus normalized radius) already outperforms the pure single-variable baseline M0, showing that the strongest operational baseline is not simply a one-variable RAR function. We therefore say that the RAR-like local baseline remains *unbeaten* by the tested static chamber augmentations, not that the RAR has been demonstrated to be sufficient.

11.5 Future programme

Three branches follow directly from the results above: **observational normalization**, formalizing $\mathcal{N}(\mathcal{F}_g^{\text{obs}}, \mathcal{O}_g)$ to separate the physical and observational layers identified in Section 9; **cluster-membership validation**, testing whether the admissibility clusters of Section 6 survive observational normalization and replicate on independent galaxy samples, rather than reflecting sample-specific structure in the primary 121-galaxy atlas; and **local ρ-based transfer models**, building on Section 10.4's identification of $\rho(\kappa)$ as the leading transfer-sensitive channel with a local, regime-conditioned, and possibly transition-based architecture rather than a static whole-curve average.

12 Conclusions

We constructed a reproducible multiscale structural-fingerprint atlas for 121 galaxy baryonic-acceleration ladders derived from resolved SPARC mass models. The atlas preserves galaxy-specific identity under unchanged preparation, exhibits channel-dependent dimensionality, separates Stable Structure and Weak Persistence regimes, resolves the Stable Structure basin into a five-family consensus admissibility structure with statistically significant differentiation in independent physical and connectivity properties, and demonstrates that persistence, fragmentation, and connectivity are distinct structural properties. It preserves recognizable physical galaxy organization while also revealing a measurable observational transformation layer. An external, held-out transfer test against RAR-like local baselines shows that these fingerprints are not adequately represented by a single static whole-galaxy predictive architecture, while the full $\rho(\kappa)$ curve emerges as the leading transfer-sensitive channel for future work. The results establish galaxy structural fingerprints as a new object of analysis and define a concrete route toward normalized, local, and regime-conditioned applications.

Data and Code Availability

The primary SPARC sample, chamber configurations, PCA tables, Gate B/C result records, and external transfer model outputs used in this manuscript are drawn from the registered UNNS_GALAXY_STRUCTURE_AND_VALIDATION_v0_3 package. The admissibility-cluster analysis of Section 6 (Ward, Gaussian-mixture, and HDBSCAN clustering; bootstrap and consensus co-assignment computation) was performed with scikit-learn, scipy, and hdbscan on the same registered PCA and chamber-signature tables. Figure- and table-generation scripts are available from the corresponding author upon request.

Acknowledgements

We thank the SPARC collaboration for the public rotation-curve database that underlies this analysis.

References

Di Cintio, A., & Lelli, F. 2016, MNRAS, 456, L127

Lelli, F., McGaugh, S. S., & Schombert, J. M. 2016, ApJL, 816, L14

Lelli, F., McGaugh, S. S., & Schombert, J. M. 2016, AJ, 152, 157 (SPARC database)

Li, P., Lelli, F., McGaugh, S. S., & Schombert, J. 2018, A&A, 615, A3

McGaugh, S. S., Lelli, F., & Schombert, J. M. 2016, PhRvL, 117, 201101

Milgrom, M. 1983, ApJ, 270, 365

Mistele, T., McGaugh, S. S., Lelli, F., Schombert, J. M., & Li, P. 2024, JCAP, 04, 020

Navarro, J. F., et al. 2017, MNRAS, 471, 1841

UNNS Structural Astrophysics Programme 2026, GAL-FINGERPRINT-I Protocol v0.1 and STRUC-I v1.0.4 / STRUC-PERC-I v2.5.0 instrument documentation, internal registered technical report

A Supplementary Figures

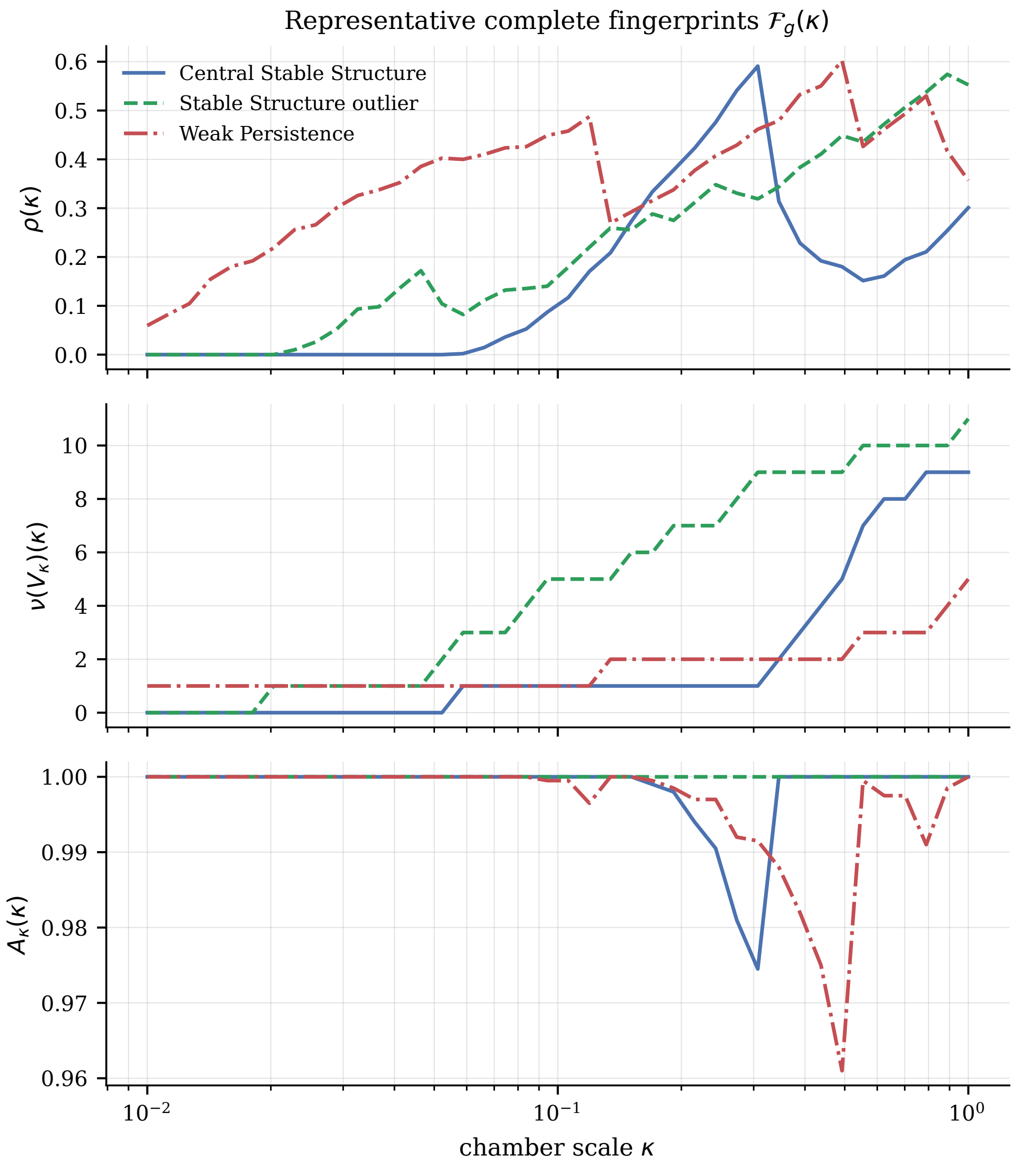

Figure 11: Representative complete fingerprints $\mathcal{F}_g(\kappa)$ for a central Stable Structure galaxy (NGC 4217), a Stable Structure outlier (UGC 11455), and a Weak Persistence galaxy (UGC 11820).

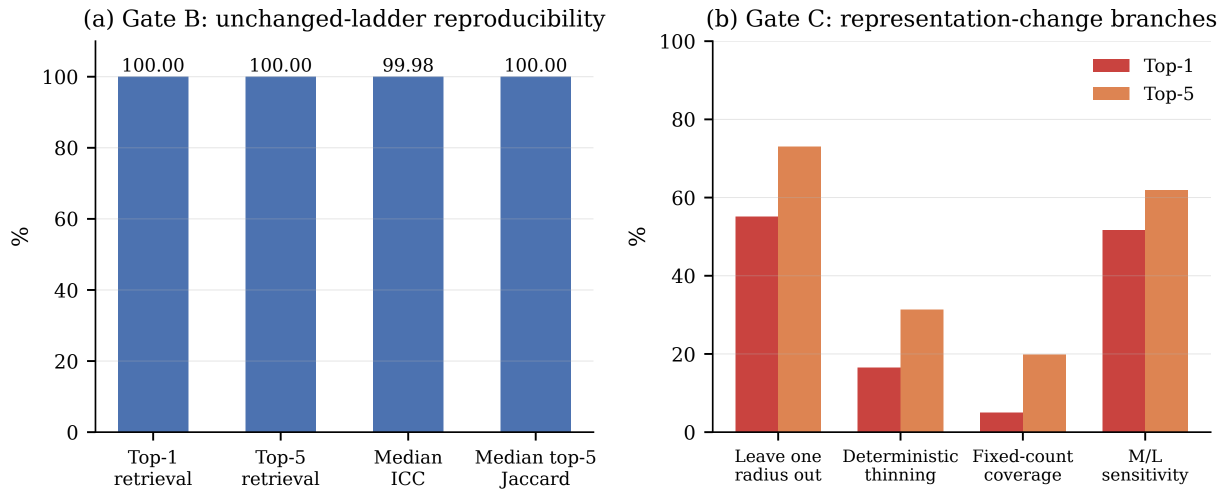

Figure 12: (a) Gate B unchanged-ladder reproducibility metrics. (b) Gate C top-1 and top-5 retrieval by representation-change branch.

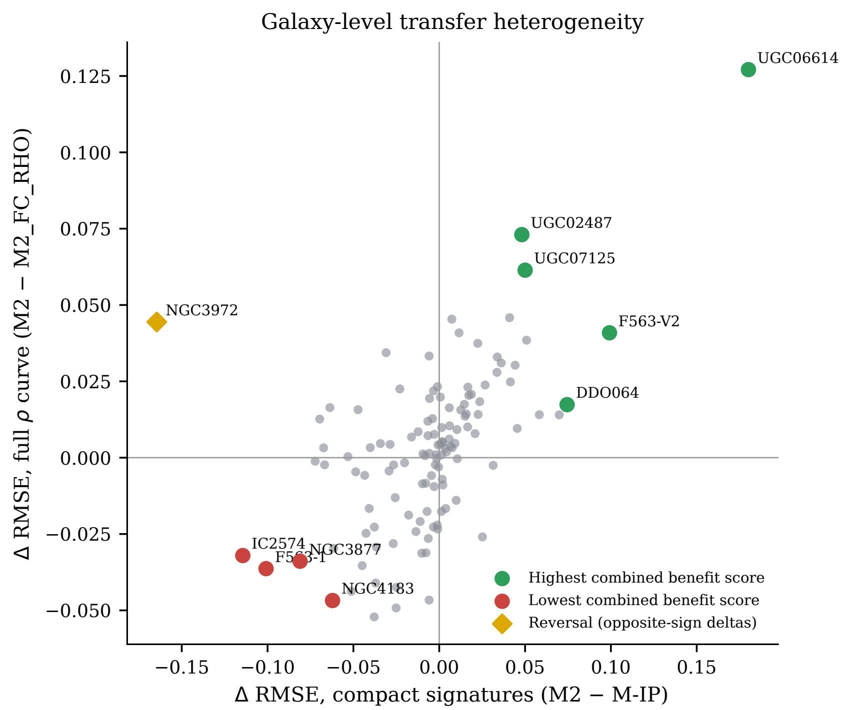

Figure 13: Galaxy-level transfer heterogeneity: per-galaxy RMSE improvement under compact signatures (M2−M-IP) versus the full ρ curve (M2−M2_FC_RHO). Green points mark the highest combined-benefit-score galaxies; red points mark the lowest combined-benefit-score galaxies. The gold diamond marks NGC 3972, the one case among these where compact signatures and the full ρ curve disagree in sign (compact signatures worsen it; the full ρ curve improves it) – a reversal, not a consistent deterioration.

Table 13: Full galaxy membership of the $k = 5$ consensus admissibility-cluster registry within Stable Structure (Table 6, Section 6.4; $n = 110$).

Cluster	Size	Members
0	65	D631-7, DDO 154, DDO 168, ESO116-G012, F563-1, F568-3, F571-8, IC 4202, NGC 0055, NGC 0100, NGC 0289, NGC 0891, NGC 1090, NGC 2683, NGC 2841, NGC 2955, NGC 3521, NGC 3726, NGC 3769, NGC 3893, NGC 4010, NGC 4013, NGC 4088, NGC 4157, NGC 4214, NGC 4559, NGC 5005, NGC 5033, NGC 5055, NGC 5371, NGC 5585, NGC 5907, NGC 6195, NGC 6503, NGC 6674, NGC 6946, NGC 7331, NGC 7814, UGC 00128, UGC 02885, UGC 02916, UGC 02953, UGC 03205, UGC 03546, UGC 04278, UGC 05005, UGC 05253, UGC 05721, UGC 05764, UGC 06614, UGC 06787, UGC 06930, UGC 07089, UGC 07125, UGC 07151, UGC 07323, UGC 07399, UGC 07524, UGC 08286, UGC 08490, UGC 09037, UGC 09133, UGC 11557, UGC 11914, UGC 12506
1	13	DDO 161, F579-V1, NGC 0300, NGC 0801, NGC 3109, UGC 01230, UGC 05750, UGC 05829, UGC 05986, UGC 06983, UGC 07603, UGC 12632, UGC 12732
2	18	DDO 064, ESO079-G014, F568-V1, F583-4, KK98-251, NGC 0024, NGC 0247, NGC 2403, NGC 2915, NGC 3877, NGC 6015, UGC 00731, UGC 01281, UGC 06446, UGC 06786, UGC 08699, UGC 11455, UGCA444
3	7	F568-1, F583-1, NGC 2976, NGC 3741, NGC 3917, NGC 4100, UGC 03580
4	7	F563-V2, NGC 2903, NGC 4183, NGC 4217, UGC 05716, UGC 06917, UGC 08550