ATLAS PAPER IV

From Coherence Surface to Extracted Boundary Geometry

Proof-of-Concept Results for Many-Source Weak-Field General Relativity

Daniel Scieszinski
Atlas Project
April 2026
danski2017@gmail.com

Abstract

The Gravitational Coherence Surface (GCS) was introduced as a weak-field boundary of source-readable gravitational distinguishability: the locus at which a chosen source signal reaches parity with a calibrated competing floor. In many-source systems, that logic predicts not an isolated shell but a structured family of balanced interfaces. We report the first proof-of-concept extraction results for that program. Using exact-source weak-field evaluations under a declared readable-field regime based on the canonical tidal proxy $Q_i \propto M_i/r_i^3$, we construct pairwise residual fields and harvest boundary-support geometry from their parity structure. Three results are emphasized. First, in extreme-ratio settings, subdominant local readable pockets persist inside globally dominant environments rather than collapsing into trivial single-source wash. Second, in dense many-source configurations, foam-like interface architectures emerge from the extraction rule itself rather than being imposed geometrically in advance. Third, under matched exact-source budgets, smooth, weak-arm, and lopsided source morphologies yield distinguishable boundary-support responses, indicating sensitivity to source arrangement beyond bulk totals alone. We interpret these results conservatively. They do not imply modified gravity, new invariants, or direct observables. They do show that many-source weak-field fields admit an operationally recoverable mesoscale boundary geometry whose morphology tracks hierarchy, anisotropy, and arrangement. Atlas is therefore positioned not as a replacement for existing gravitational solvers, but as an extraction instrument for derived readability structure in many-source gravitational environments.

Keywords: weak-field general relativity; many-source gravity; gravitational coherence surface; parity extraction; mesoscale geometry; morphology-sensitive boundary structure; proof-of-concept methods

1. Introduction

Gravitational theory already provides a complete local language for weak-field source evaluation through the metric, curvature, and geodesic-deviation structure of standard general relativity. What it does not automatically provide is a mesoscale descriptive language for how readable source structure is partitioned, shared, or handed off across a many-source environment. The original Gravitational Coherence Surface formalism addressed one part of that gap by defining a boundary of source-readable distinguishability: the locus at which a chosen source signal reaches parity with a calibrated competing floor. In that formulation, the GCS was not a force cutoff, not a stability boundary, and not a modification of gravity. It was a derived bookkeeping geometry for where source-specific gravitational readability ceases to stand out cleanly against its environment.

Once several comparable sources are allowed to compete on the same weak-field slice, however, the single-shell picture is no longer sufficient. The many-source extension of the coherence-boundary logic predicts not an isolated enclosing surface but a structured family of parity supports: balanced interfaces, exposed faces, seam-like intersections, and related support objects arising from source-readable competition. In earlier Atlas work, this many-source picture was embedded within a broader datum-centered witness architecture in which root datums, node-specific witness structures, overlap comparisons, and canonical parity extraction form a single operational pipeline. Within that pipeline, the shared parity network was assigned a disciplined role: not the whole ontology of the solver, but the first stable extracted artifact family within a broader comparison engine.

A related concern, emphasized in subsequent Atlas work, is that many-source gravitational fields may contain a neglected mesoscale residue layer between pointwise local evaluation and coarse averaging. In that middle regime, structured competition may remain geometrically legible before conventional smoothing or bulk summaries erase it. If so, one should be able to extract support geometry that is sensitive not only to total source content, but also to hierarchy, anisotropy, and morphology.

This paper reports the first proof-of-concept extraction results for that program. Working entirely within standard weak-field general relativity, we adopt an exact-source readable regime based on the canonical tidal proxy shown below, construct source-resolved comparison fields, and harvest boundary-support geometry from their parity structure. The purpose is not to claim a new invariant, a new force law, or an observational detection. The purpose is narrower and more foundational: to establish that the coherence-boundary program has advanced from formal proposal to operational extraction.

$$Q_i(x) \propto M_i / r_i^3$$

The results are interpreted conservatively. They do not establish that every readable operator yields the same geometry, that the extracted structures are physical observables, or that the resulting artifacts are fundamental invariants of the field. They do establish that many-source weak-field gravity admits an operationally recoverable derived boundary geometry. Atlas is therefore best understood, at this stage, not as a replacement for existing gravitational solvers, but as an extraction instrument for readable mesoscale structure in exact-source many-body environments.

2. Paper position within the Atlas sequence

This paper occupies a specific place in the Atlas sequence. Paper I introduced the Gravitational Coherence Surface as a boundary of gravitational distinguishability within standard general relativity, with explicit separation between signal readability and dynamical binding. Paper II generalized that logic to many-source systems and repositioned Atlas as a datum-centered curvature witness engine whose first canonical artifact family is the shared parity network. Paper III then argued that the overlap layer between pointwise evaluation and coarse averaging may contain a neglected mesoscale legibility stratum, in which interface geometry, low-margin competition, and seam structure become readable before averaging erases them.

The present paper serves a different purpose. It is the first results paper. Its task is to show that the Atlas program has moved beyond formal proposal and into operational extraction. The central question here is no longer merely whether the coherence-boundary idea is mathematically definable. The question is whether a declared readable comparison rule, applied to exact-source weak-field fields, actually recovers reproducible many-source boundary geometry. The answer given by the proof-of-concept runs reported here is yes, in a disciplined first sense. Atlas can recover boundary-support structures that deform under hierarchy, anisotropy, and morphology, and it can do so in ways that are visibly tied to source arrangement rather than to arbitrary rendering choices alone.

3. Declared regime and discipline of claim

The present paper remains intentionally narrow in regime. All results are taken within standard weak-field general relativity. No new field equations are introduced. No modification of gravity is proposed. The extraction is based on exact-source evaluations under a declared readable rule, and the canonical readable branch is the weak-field tidal proxy fixed as the first Atlas milestone language because it is analytically transparent, benchmark-friendly, and sufficient to stabilize the support-extraction pipeline before richer tensorial branches are promoted.

$$Q_i(x) \propto M_i / r_i^3$$

Given a set of sources S_i, the first comparison objects are the pairwise residual fields shown below. Their zero sets define candidate parity supports. In the simplest two-source point-mass specialization, this recovers midpoint sheets in the equal-mass case and mass-bent sheets in the unequal-mass case. Atlas is therefore support-first rather than mesh-first. The primary objects are harvested supports, retained boundary points, and certified low-margin structures, with rendered surfaces and plots functioning as downstream expression layers rather than as sources of truth.

$$\Psi_{ij}(x) = Q_i(x) - Q_j(x)$$
$$M_i / r_i^3 = M_j / r_j^3$$

The discipline of claim is correspondingly conservative. We do not claim that every readable operator will produce the same boundary geometry. We do not claim that the extracted supports are physical observables in their present form. We do not claim new invariants. We do claim that, once a readable operator is fixed, many-source weak-field fields admit a derived boundary geometry that can be operationally extracted and compared across configurations. That claim is weaker than ontology and stronger than metaphor. It is the correct claim for a first evidence paper.

4. Operational extraction method

The operational method used in these proof-of-concept runs follows the canonical Atlas extraction logic. A source configuration is declared on a chosen slice. Exact-source readable quantities are evaluated directly from that configuration under the fixed weak-field proxy $Q_i \propto M_i/r_i^3$. From those evaluations, pairwise residual fields are formed. Boundary-support geometry is then harvested from parity conditions, low-margin regions, or dense support clusters, depending on the local interrogation mode. The retained products are not assumed surfaces in advance. They are extracted supports that may later be rendered as point clouds, slices, or reconstructed sheets for inspection.

The proof-of-concept images reported here come from three operational contexts. The first is an extreme-ratio local-pocket test, designed to determine whether a light-source-readable region can remain legible inside a field dominated globally by a much stronger competitor. The second is a dense many-source isotropic boundary projection, designed to test whether boundary-support geometry emerges at all as a harvestable artifact in a complex source environment. The third is a matched-budget morphology comparison, in which smooth, weak-arm, and lopsided source arrangements are interrogated under the same exact-source budget to test whether the extracted structure responds to morphology rather than merely to total source content. These three contexts form a minimal but revealing first benchmark triad: hierarchy, emergence, and morphology.

5. Results

5.1. Extreme-ratio environments preserve subdominant local readable structure

The first proof-of-concept test examines whether Atlas can recover nontrivial readable structure around a local source embedded within a much stronger surrounding field. If the method simply follows the globally dominant contribution everywhere, then extreme-ratio systems should collapse into trivial single-source outputs with little or no retained local structure. If, however, readable competition remains locally meaningful under strong hierarchy, then one expects a finite extracted pocket around the subordinate source, deformed by the surrounding field rather than erased by it.

That second behavior is what the present extraction shows. Figure 1 exhibits a bounded local support region centered on the anchor source but distorted in the direction of the dominant competitor. The recovered geometry is therefore neither isotropic nor numerically diffuse. It is a shaped local parity-support pocket, compressed and displaced by the larger environment but still operationally present. This is precisely the regime in which a boundary-extraction method must prove that it can preserve subdominant legibility rather than merely report the largest contribution everywhere.

The result is significant for two reasons. First, it shows that the Atlas extraction is not restricted to near-symmetric benchmark scenes. It retains readable local structure even when the global field is strongly hierarchical. Second, it supports a broader methodological point already emphasized in the internal architecture work: extreme hierarchy should be treated as a scale-transition problem rather than as a solver failure, and locally coherent pockets should not be erased simply because coarse global interrogation under-resolves them.

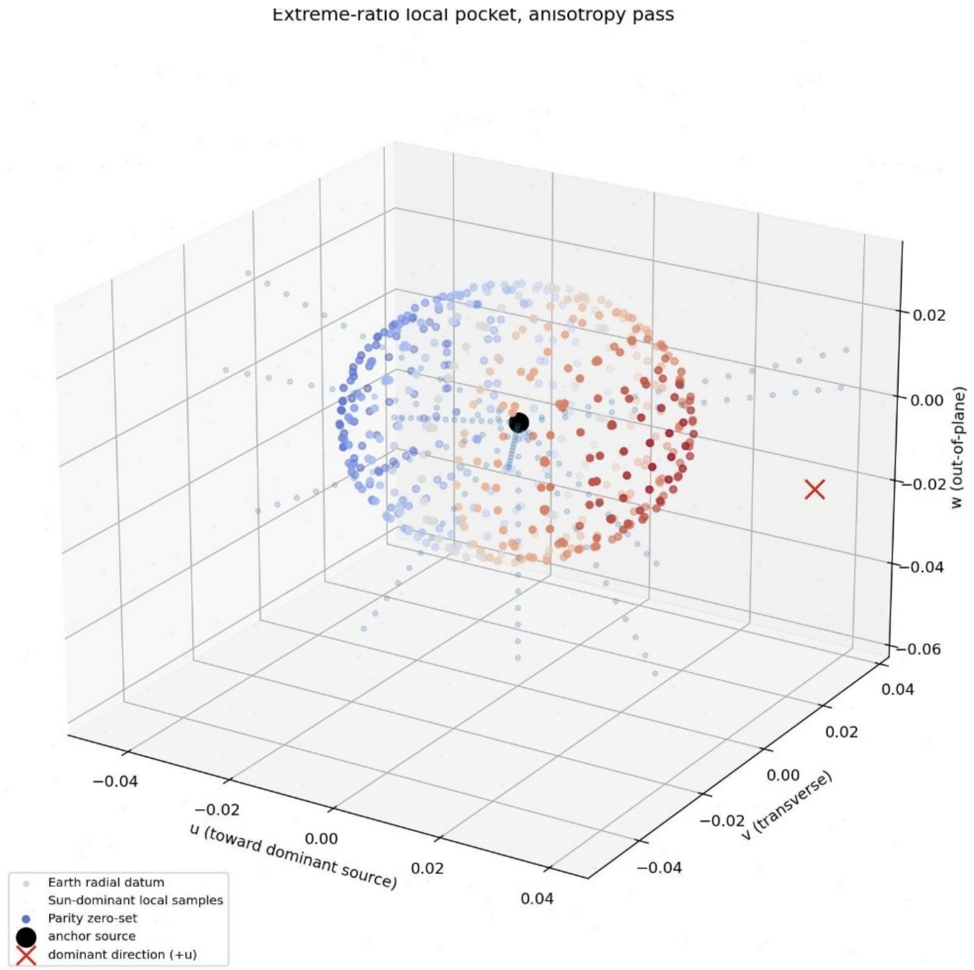

Figure 1. Extreme-ratio local-pocket extraction under anisotropic competition. Boundary-support points recovered around a local anchor source in the presence of a much stronger dominant competitor. The extracted pocket remains finite and visibly anisotropic rather than collapsing into a trivial single-source wash. Color encodes local readable dominance relative to the dominant-source direction. This figure is included as proof of concept that Atlas can preserve subdominant local readable structure within a globally asymmetric weak-field environment.

5.2. Many-source readable competition condenses into harvestable boundary architecture

The second result addresses a more basic question: does the many-source boundary logic yield a dense extractable architecture at all when source competition becomes complex? Paper II argued that, once the GCS logic is lifted into genuine many-source settings, pairwise readable balance should generate balanced interfaces, exposed faces, and seam-like structures rather than a single enclosing shell. The present extraction provides the first direct proof-of-concept demonstration of that transition in harvested support form.

Figure 2 shows a high-density many-source boundary harvest projected into a neutral isotropic slice. The resulting support pattern has a visibly foam-like morphology. This description is deliberately geometric rather than ontological. We do not claim a literal physical foam or any new microscopic structure of spacetime. What the figure demonstrates is narrower and more important: under exact-source weak-field evaluation and a declared readable comparison rule, interface-bearing support geometry can emerge as a product of the extraction itself. The partition-like architecture is not imposed as a prior tessellation. It is recovered from readable competition.

This is the point at which the Atlas program moves from conceptual picture to instrument behavior. Theoretical language such as shared parity network or balanced interface family is meaningful only if it can be made to produce stable, inspectable structure under explicit evaluation rules. Figure 2 shows that such structure is in fact harvestable.

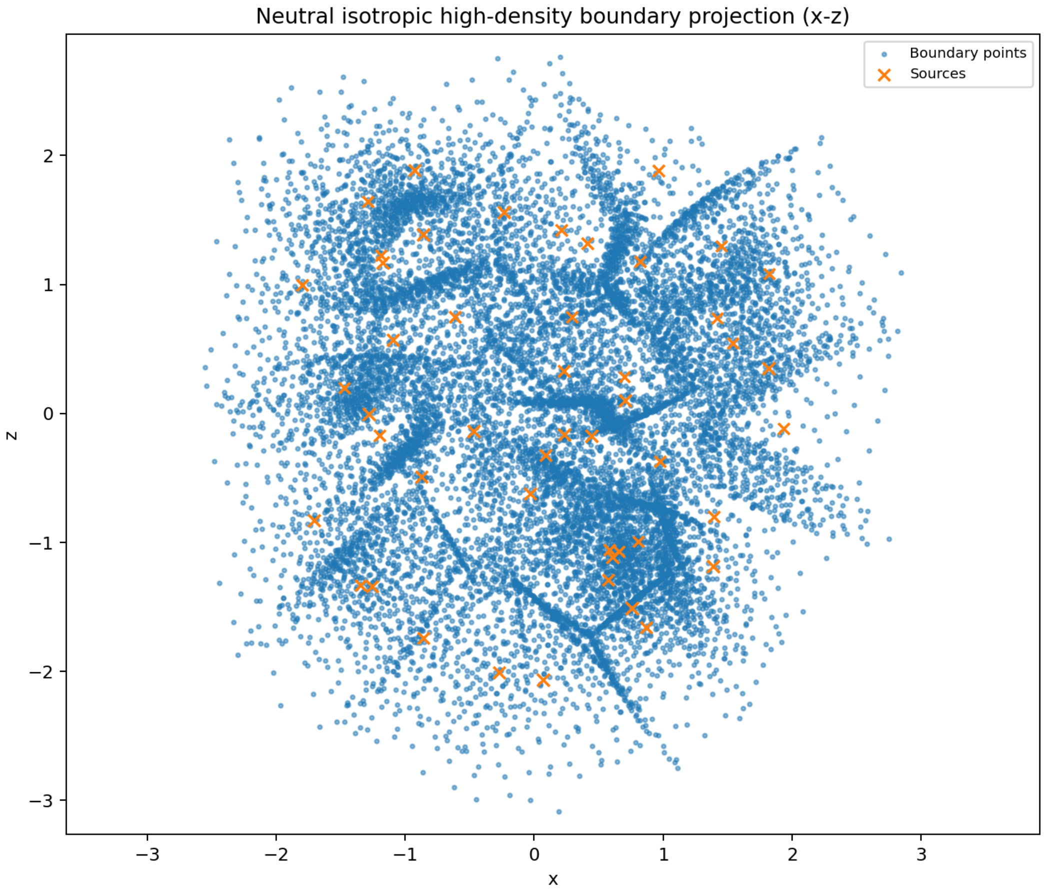

Figure 2. Dense many-source boundary harvest in a neutral isotropic projection. High-density boundary-support points extracted from a many-source weak-field configuration and projected into the x-z plane. Orange markers denote source locations; blue points denote retained boundary-support samples. The resulting architecture exhibits a foam-like partition morphology that emerges from the exact-source readable comparison rule itself rather than from any imposed tessellation prior.

5.3. Matched-budget source ensembles yield morphology-sensitive boundary responses

The third result is the most important for physical interpretation because it asks whether the extracted boundary geometry responds to arrangement rather than merely to total source content. Three ensembles with the same exact-source budget were interrogated under identical extraction logic: a smooth distribution, a weak-arm distribution, and a lopsided distribution. If the Atlas outputs were largely insensitive to morphology, then these three cases should produce approximately degenerate support slices after normalization by source count or total budget. Instead, they produce visibly distinct boundary architectures.

Figure 3 shows that the extracted support patterns reorganize across the three ensembles. The smooth case yields a relatively even and compact boundary field. The weak-arm case develops directionally biased extension and structured asymmetry. The lopsided case shows stronger displacement, imbalance, and concentration into preferential sectors. These differences are not incidental. They indicate that the extraction responds to how source content is distributed in space, not merely to how much source content is present.

This is the first proof-of-concept evidence that the Atlas boundary program may be sensitive to morphology in exactly the way Paper III suggested: not as a replacement for bulk descriptors, but as a middle-layer residue geometry that remains legible before coarse averaging compresses the field into broader summaries. What Figure 3 establishes is not yet a morphology invariant or a quantitative classifier. It establishes non-degenerate response under matched budget, and that is already a strong result.

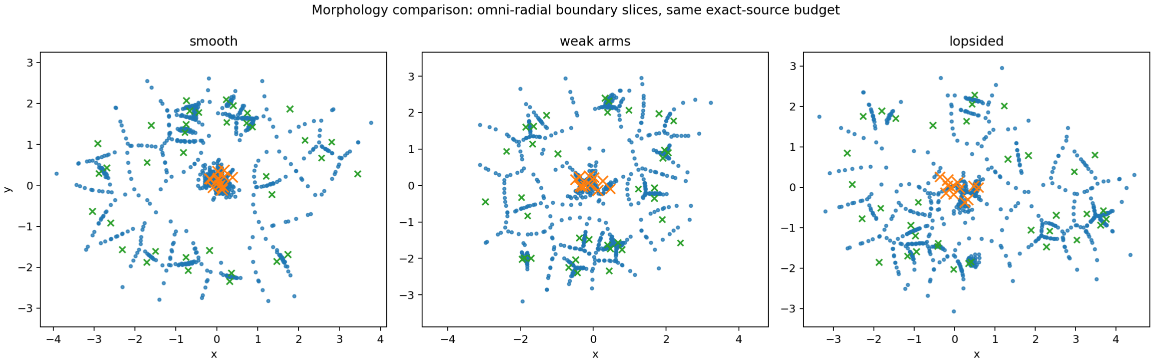

Figure 3. Morphology sensitivity under matched exact-source budget. Boundary-support slices extracted for three source ensembles with the same exact-source budget but different spatial organization: smooth, weak arms, and lopsided. Blue points show retained boundary supports; source markers show the corresponding underlying source layouts. The non-degenerate differences among the three extracted patterns indicate that the Atlas boundary geometry responds to source morphology rather than only to bulk source count or total budget.

5.4. Cross-result synthesis

Taken together, Figures 1 through 3 reveal a coherent pattern. The first shows that readable support can survive local hierarchy. The second shows that many-source competition can condense into dense harvestable interface architecture. The third shows that the resulting architecture is not generic, but deforms in response to source morphology. These three behaviors correspond to three minimal criteria for an extraction instrument of the kind Atlas proposes.

The first criterion is persistence under asymmetry. The instrument must not fail as soon as a dominant source enters the scene. The second is emergence under complexity. The instrument must recover structure in multi-source settings rather than only in analytically simple benchmarks. The third is sensitivity to arrangement. The instrument must distinguish source layouts that share comparable budgets but differ in geometry. The present proof-of-concept results satisfy all three criteria in a first operational sense.

6. Discussion

6.1. What has been established

The proof-of-concept results reported here establish a threshold claim: many-source weak-field gravitational fields admit an operationally recoverable boundary-support geometry under a declared readable comparison rule. That statement is narrower than ontology and broader than visualization. It does not say that the extracted structures are fundamental objects of the theory in the same sense as the metric or curvature tensor. It does say that, once a source-readable regime is fixed, Atlas can recover structured support artifacts whose behavior is systematic enough to benchmark. That is the correct level of claim for a first evidence paper.

Three specific properties have been demonstrated. First, readable support can persist locally inside globally dominant environments, indicating that extraction does not simply collapse to the strongest source everywhere. Second, dense many-source readable competition can yield a harvestable interface architecture rather than only isolated parity traces. Third, that architecture is sensitive to source arrangement under matched budget, which means the method is responding to morphology rather than only to total content. Taken together, these are the first signs that the coherence-boundary program is recovering a genuine mesoscale descriptive layer in many-source weak-field fields.

6.2. Why this matters physically

The physical importance of these results lies in description, not in revised dynamics. Standard general relativity already determines the local field. What Atlas adds is a way to organize that field into readable support structure: where one source dominates, where readable competition becomes shared, and how those transitions deform under hierarchy and morphology. This is useful precisely because many-source systems often occupy an uncomfortable descriptive middle ground. Pointwise evaluation is exact but atomized. Coarse summary is compact but can erase relational structure. Questions of averaging and coarse-graining in relativistic settings are longstanding; the present paper is best read as an extraction-layer contribution alongside, not against, that broader literature. Atlas is aimed at the middle layer between those extremes.

The present results suggest that this middle layer is not empty. In particular, the morphology-sensitive response of Figure 3 implies that the extracted support geometry retains information about spatial arrangement that is not obviously reducible to a bulk scalar summary. Atlas is therefore not intended as a replacement for mature simulation infrastructures, but as a derived-geometry extraction layer that could sit beside or downstream of them. That does not yet imply a new observable. It does imply a new comparative language. Systems with similar total budget but different organization need not be equivalent in their readable boundary structure.

6.3. Relation to the earlier Atlas papers

These results also clarify the internal logic of the Atlas paper sequence. Paper I established that gravitational readability can be bounded without being confused with dynamical binding. Paper II argued that many-source competition should generalize the single-shell picture into a network of balanced supports inside a datum-centered witness architecture. Paper III proposed that the overlap layer might contain persistent residue geometry worth extracting before averaging washes it out. The present paper does not replace any of those claims. It operationalizes them. It shows that the program can now produce support artifacts substantial enough to be treated as results rather than merely as motivating concepts.

6.4. What remains open

The current results are intentionally preliminary. Several major questions remain open. The first concerns persistence. A proof-of-concept extraction is not enough by itself. The retained supports must now be tested under perturbations of source layout, sampling density, domain size, and readable branch to determine which structures are stable and which are merely provisional.

The second concerns certification. Atlas has already emphasized a support-first discipline in which rendered surfaces remain downstream of certified support rather than becoming truth by visual charisma. That discipline now has to mature into explicit certification logic: which support objects count as primary, which are secondary, and what persistence gates must be passed before an artifact is promoted.

The third concerns branch comparison. The canonical readable proxy $Q_i \propto M_i/r_i^3$ remains the correct first milestone language because it is analytically transparent and operationally stable. But one eventual question is how the extracted geometry changes when richer readable branches are activated, including directional tensor projections or norm-based weak-field tidal measures. The present paper does not answer that question. It establishes the baseline against which that question can later be asked cleanly.

The fourth concerns physical correlation. A derived boundary geometry becomes more scientifically significant if it sharpens downstream tasks such as geodesic comparison, environment-sensitive inference, or morphology-sensitive classification. The present paper shows that there is something to correlate. It does not yet demonstrate those correlations.

6.5. Limits of interpretation

The strongest interpretation risk at this stage is over-promotion. A foam-like boundary harvest is not yet evidence of a new physical medium. A morphology-sensitive support field is not yet a new invariant. A locally persistent readable pocket in an extreme-ratio scene is not yet a direct observational claim. The present paper should therefore be read as a disciplined extraction result, not as an ontological verdict.

That distinction matters because Atlas is most credible when it preserves the hierarchy between field physics and derived geometry. The local field remains primary. The extracted support structure is secondary and induced. Its scientific value must be earned through stability, comparative usefulness, and eventual correlation with other physically meaningful diagnostics.

6.6. Immediate next phase

The immediate next phase is benchmark hardening rather than conceptual expansion. The logical program is already visible. Begin with exact symmetric and near-symmetric benchmarks. Continue through unequal-mass deformation and shape-perturbed families. Compare extracted supports against the omni-radial null model and require any more elaborate refinement logic to justify itself by recovering new structure, cleaner lineage, or materially better local fidelity. Keep support certification ahead of mesh expression. Treat richer readable branches as validation lanes rather than coequal architecture until the canonical branch is fully stabilized.

A useful way to summarize the present state is this: Atlas has not yet proven a new ontology of gravitational structure, but it has demonstrated a new instrument behavior. It can extract a boundary-support layer from many-source weak-field fields, and that layer responds to hierarchy, anisotropy, and morphology. That is enough to move the coherence-surface program out of the realm of pure formal suggestion and into the domain of benchmarkable physics.

7. Conclusion

The Gravitational Coherence Surface was introduced as a boundary of source-readable gravitational distinguishability within standard weak-field general relativity. In many-source systems, that logic predicts not a single enclosing shell but a structured family of balanced supports arising from readable competition. The present paper has reported the first proof-of-concept extraction results for that broader program. Under a declared exact-source readable regime based on the canonical tidal proxy $Q_i \propto M_i/r_i^3$, Atlas recovers boundary-support geometry that is locally persistent under hierarchy, globally structured under dense many-source competition, and visibly responsive to source morphology under matched exact-source budget.

These results do not establish a new invariant, a new dynamical law, or a new observational class. They do establish something narrower and foundational: the coherence-boundary program has become operational. Atlas can now extract a derived mesoscale geometry from many-source weak-field fields rather than merely describe one in principle. That geometry is best understood, at this stage, as a boundary-support layer induced by readable competition within standard general relativity, not as a replacement for the underlying field description.

The immediate scientific value of this result is methodological and comparative. It provides a new descriptive instrument for asking where source readability persists, where it becomes shared, and how those transitions deform under changes in hierarchy, anisotropy, and morphology. In that sense, this paper marks a threshold rather than a culmination. Paper I defined the readability boundary. Paper II positioned the many-source witness architecture. Paper III argued for the relevance of the mesoscale overlap layer. Paper IV shows that the instrument can now recover that layer in proof-of-concept form. The next phase is to place these extracted structures under stricter benchmark discipline and determine which of them deserve promotion from interesting artifacts to stable comparative geometry.

Appendix A. Minimal methods and conventions

All proof-of-concept figures reported here were generated from Atlas runs operating in the weak-field exact-source regime. The declared readable field was the canonical scalar tidal proxy $Q_i \propto M_i/r_i^3$. Source-resolved competition was evaluated directly from the declared source ensemble, not from a coarse-grained density field or imposed tessellation prior.

The operational pipeline may be summarized as follows: source declaration, readable-field evaluation, pairwise or local competition assessment, boundary-support harvesting, and downstream figure rendering. The rendered figures should therefore be interpreted as visual presentations of extracted support artifacts rather than as the extraction itself. In Atlas terminology, support remains primary and mesh or point-cloud presentation remains downstream expression.

The three figures serve different diagnostic roles. Figure 1 probes local readable persistence under strong hierarchy. Figure 2 probes whether dense many-source readable competition yields extractable interface-bearing structure in a neutral isotropic view. Figure 3 probes whether matched-budget ensembles differing in spatial organization yield non-degenerate support responses. Together, they provide a minimal first benchmark triad for the present paper: hierarchy, emergence, and morphology.

Appendix B. Figure notes and interpretation discipline

The phrase foam-like is used in this paper as a visual descriptor of the recovered partition morphology. It is not intended as a literal ontological claim about spacetime microstructure. Similarly, the terms pocket, boundary-support geometry, and interface architecture should be read as names for induced readable structures inside a declared weak-field comparison regime.

The present paper is intentionally silent on strong claims of uniqueness or observational equivalence. The appropriate scientific role of these figures is proof of concept: they demonstrate that Atlas can recover systematic structure from many-source readable competition. They do not by themselves settle how branch-stable, invariant, or directly observable that structure may become under later benchmark work.

References

Buchert, T. (2000). On average properties of inhomogeneous cosmologies. General Relativity and Gravitation, 32, 105-125.

Löffler, F., et al. (2012). The Einstein Toolkit: A community computational infrastructure for relativistic astrophysics. Classical and Quantum Gravity, 29, 115001.

Souami, D., Cresson, J., Biernacki, C., & Pierret, F. (2020). On the local and global properties of gravitational spheres of influence. Monthly Notices of the Royal Astronomical Society, 496, 4287-4300.

Zalaletdinov, R. M. (2008). The averaging problem in cosmology and macroscopic gravity. International Journal of Modern Physics A, 23, 1173-1181.

Scieszinski, D. (2026). The Gravitational Coherence Surface: Signal-to-Noise Limits on Gravitational Identity. Atlas Paper I / Faculty Review Draft.

Scieszinski, D. (2026). Shared Parity Networks within a Datum-Centered Curvature Witness Engine. Atlas Paper II, revised architecture draft.

Scieszinski, D. (2026). Mesoscale Residues and the Derived Geometry of Multipole Tidal Composition. Atlas Paper III.

Atlas Project. (2026). The Gravitational Coherence Surface and Shared Parity Network Framework: A Canonical Postulate with Interpretive Extension. Internal continuity document.

Atlas Project. (2026). Shared Parity Network Breakthrough and Extraction-Pipeline Lock-In. Session note.

Atlas Solver Project. (2026). Project Tracker. Consolidated working tracker, updated April 13, 2026.