Z–ZOA: A Probabilistic 3D Atlas of the Zone of Avoidance

A framework for reconstructing hidden large-scale structure and testing its dynamical consequences

Zhane Grey
Independent Researcher
Z–OA Project

14 September 2026

Abstract

The Galactic Zone of Avoidance (ZoA) is not an empty region of the nearby Universe. It is an observationally heterogeneous region in which Galactic dust, stellar crowding, radio-frequency interference, continuum emission, and wavelength-dependent galaxy-selection effects suppress or distort the apparent distribution of extragalactic structure. Dedicated H I surveys have already revealed galaxies and walls hidden from optical catalogues, including a prominent continuation of the Great Attractor wall in the inner Galactic plane and multiple structures associated with the Vela Supercluster. The next scientific opportunity is therefore not to "discover" the ZoA, but to unify its fragmented observations into an uncertainty-aware, three-dimensional reconstruction.

This paper proposes Z–ZOA, a reproducible framework that combines blind H I catalogues, near-infrared catalogues, optical and spectroscopic redshifts, peculiar-velocity measurements, survey footprints, and explicit selection functions. The framework represents the observed galaxy field and the unobserved or incompletely observed field as separate probabilistic layers. It then propagates catalogue uncertainties into density reconstruction, filament and wall topology, velocity-field inference, watershed basin classification, and the predicted gravitational acceleration of the Local Group. Two staged analyses are proposed: a nearby run with $0 < cz < 12{,}000$ km s^{-1} and a deeper run extending to 25,000 km s^{-1}.

The principal test is whether filling the ZoA with radio and near-infrared constraints changes inferred large-scale structure and local cosmic-flow basins in a statistically meaningful way. The project is designed as a methods and data-integration study; it reports no new reconstruction results until the proposed master catalogue and posterior analyses are executed. Its intended contribution is a transparent bridge between modern ZoA surveys and Cosmicflows-style dynamical cosmography.

Keywords: large-scale structure of Universe; Zone of Avoidance; neutral hydrogen; galaxy catalogues; peculiar velocities; cosmic flows; Great Attractor; Vela Supercluster; Bayesian inference

1 Introduction

Maps of the nearby Universe are shaped as much by the observing instrument as by the cosmic web itself. The Milky Way occupies a foreground screen across the sky. At low Galactic latitude, dust extinction reduces optical detectability, foreground stars create severe source confusion, and bright Galactic radio emission complicates even radio observations. The resulting band is conventionally called the Zone of Avoidance. It is a limitation on the observed galaxy field, not a physical prediction that the matter density is near zero.

This distinction matters dynamically. The Great Attractor and Norma region, the Local Void, the Hydra–Centaurus complex, the Vela Supercluster, and connected walls and filaments lie in directions intersected by the Galactic plane. The missing or biased part of the galaxy distribution therefore overlaps structures that may contribute to the Local Group's peculiar motion and to the inferred shape of local velocity-field basins. A map that simply blanks the ZoA can hide both overdensities and underdensities while leaving downstream dynamical calculations apparently well defined.

The observational situation has improved substantially. The Parkes H I Zone of Avoidance Survey (HIZOA) produced a blind southern-plane sample of 883 galaxies to a recessional velocity of approximately 12,000 $\mathrm{km\,s^{-1}}$ (Staveley-Smith et al., 2016). MeerKAT observations have since reached much deeper and finer angular resolution in selected strips. The first reported inner-plane Great Attractor analysis identified 477 H I galaxy candidates and found 214 detections concentrated in the Great Attractor velocity interval (Steyn et al., 2024). In the Vela region, the Vela–SMGPS and Vela–H I programmes together have revealed more than 1,500 newly identified H I-rich galaxies and evidence for two prominent walls with possible crossings of the Galactic plane (Rajohnson et al., 2024a, 2026).

At the same time, near-infrared catalogues provide a complementary, but not unbiased, view. The 2MASS ZoA catalogue contains 3,763 bright galaxies at low Galactic latitude and was explicitly constructed for large-scale-structure and flow-field work (Schröder et al., 2019). Cosmicflows-4 provides independent distance information for 55,877 galaxies assembled into 38,065 groups, enabling peculiar-velocity constraints beyond redshift space alone (Tully et al., 2023). WALLABY adds a growing southern H I survey layer; by the February 2026 collaboration update, approximately 8,000 H I detections were available to the team for science analysis (WALLABY Collaboration, 2026).

These resources motivate a focused question:

> How much do observational incompleteness and newly detected ZoA galaxies change the inferred three-dimensional structure and dynamics of the nearby cosmic web?

Z–ZOA is proposed to answer that question by constructing a probabilistic atlas rather than another uncalibrated scatter plot. Its central mapping is

$$\boxed{\begin{array}{c}\text{observed galaxies} + \text{ZoA surveys} + \text{peculiar velocities} \\ + \text{ selection functions} \longrightarrow \text{probabilistic 3D cosmic map}\end{array}} \tag{1}$$

The paper has four goals. First, it defines a unified data model that preserves survey provenance and measurement uncertainty. Second, it specifies a selection-function model that distinguishes an observationally blank voxel from a physically underdense voxel. Third, it proposes tests of wall and filament connectivity in the Great Attractor and Vela regions. Fourth, it quantifies the effect of ZoA completion on reconstructed velocity basins and the predicted Local Group acceleration. The analysis is deliberately framed as a proposed study: the hypotheses and evaluation metrics are specified here, while numerical results are left to the subsequent data release and inference run.

2 Scientific context

2.1 The ZoA as a selection problem

The ZoA is a wavelength-dependent selection problem. Optical surveys are strongly affected by extinction and foreground stellar density. Near-infrared surveys reduce the impact of dust but

remain limited by stellar confusion, source blending, extinction corrections, and the incompleteness of redshift follow-up. Blind H I surveys are able to detect gas-rich galaxies through much of the dust screen and return a redshift from the 21-cm line, but they introduce a population bias toward galaxies with detectable neutral gas. Massive gas-poor ellipticals and environmentally stripped systems may therefore be underrepresented in H I-selected samples.

The consequence is that there is no single universal "ZoA catalogue." Different surveys sample different latent populations with different sensitivity, angular resolution, velocity resolution, footprints, and failure modes. A scientifically useful atlas must retain those differences rather than collapse them into a binary observed/unobserved mask.

2.2 What recent surveys already establish

The proposed project builds on established observations rather than claiming priority over them.

- HIZOA provides a deep, blind Parkes reference sample across the southern ZoA, with 883 galaxies and a nominal velocity reach of 12,000 $\mathrm{km\,s^{-1}}$ (Staveley-Smith et al., 2016).
- MeerKAT SMGPS observations toward the Great Attractor cover a narrow inner-plane strip near $302^\circ \lesssim \ell \lesssim 332^\circ$. The published sample contains 477 candidates, including a 214-object overdensity in the Great Attractor distance range (Steyn et al., 2024).
- The Vela–SMGPS survey found 843 heavily obscured galaxies in a 90 $\mathrm{deg^2}$ strip, while the adjoining Vela–H I survey catalogued 719 galaxies over approximately 242 $\mathrm{deg^2}$ (Rajohnson et al., 2024a). The combined analysis traces structures out to roughly 25,000 $\mathrm{km\,s^{-1}}$ and reports a diagonal wall and two prominent Vela-associated walls.
- The 2MASS ZoA catalogue contributes a homogeneous near-infrared sample of 3,763 galaxies at $|b| \leq 10^\circ$, with explicit analysis of extinction, stellar-density limits, and flow-field relevance (Schröder et al., 2019).
- WALLABY supplies a broad, evolving southern H I layer. Its survey design targets the local Universe and includes regions containing major nearby overdensities and voids; the collaboration's February 2026 update reports approximately 8,000 team-accessible H I detections (WALLABY Collaboration, 2021, 2026).

The scientific gap is therefore not the absence of observations. It is the absence of a common probabilistic representation in which heterogeneous detections, non-detections, selection effects, distance information, and dynamical consequences can be evaluated together.

2.3 Connection to cosmic-flow cosmography

Cosmicflows-4 demonstrates the value of combining redshifts with independent distance indicators. For a galaxy with measured distance d and observed recession velocity cz, a first-order line-of-sight peculiar velocity can be written as

$$u_{\parallel} \simeq cz - H_0 d, \tag{2}$$

where the sign convention and frame transformations must be held fixed throughout the analysis. Equation 2 is not itself a complete dynamical reconstruction: distance errors can be large, grouping changes the effective uncertainty, and redshift-space distortions complicate the relationship between tracer positions and the underlying density field. It is nevertheless the observational bridge between the hidden galaxy distribution and the velocity field.

The Laniakea watershed analysis showed how reconstructed velocity streamlines can be used to define basins of attraction in the nearby Universe (Tully et al., 2014). Z–ZOA adopts the same broad conceptual link while testing how stable such basin assignments are when the Galactic-plane data are changed. The relevant result is not whether a named supercluster boundary is visually attractive, but whether a boundary persists under survey-specific selection functions, catalogue perturbations, and alternative density-to-velocity models.

3 Research questions and falsifiable hypotheses

Z–ZOA is organized around four questions.

3.1 Great Attractor continuity

Does the Great Attractor wall form a statistically supported continuous or near-continuous structure across the inner plane? The test must distinguish a connected wall from a chance alignment in redshift space. The primary evidence will be the posterior probability of a density ridge connecting the higher-latitude structures to the MeerKAT inner-plane detections, after accounting for the survey footprint and detection threshold.

3.2 Vela morphology

Is Vela better described as a single wall, two walls, a connected knot, or a more branched web? The Vela data already show multiple structures and possible plane crossings (Rajohnson et al., 2024a, 2026). Z–ZOA will treat these descriptions as competing topological models and compare them using posterior predictive checks rather than relying on a single visual interpretation.

3.3 Basin sensitivity

Does adding ZoA constraints alter the Laniakea–Shapley watershed or other local flow-basin assignments? Let $B_0(\mathbf{x})$ be the basin label from a baseline reconstruction and $B_{\mathrm{ZoA}}(\mathbf{x})$ the label after adding the ZoA data. The proposed comparison is

$$\Delta B(\mathbf{x}) = B_{\mathrm{ZoA}}(\mathbf{x}) - B_0(\mathbf{x}), \tag{3}$$

under a label-aligned representation. Because basin labels are categorical, the principal statistics will be boundary displacement, transition matrices, Jaccard overlap, and posterior probabilities of reassignment rather than an arbitrary numeric difference.

3.4 Local Group acceleration

Does the additional structure reduce the discrepancy between the gravitational acceleration predicted by the reconstructed density field and the observed Local Group motion? The angular comparison is

$$\theta = \cos^{-1}\left(\frac{\mathbf{g}_{\mathrm{rec}} \cdot \mathbf{v}_{\mathrm{LG}}}{|\mathbf{g}_{\mathrm{rec}}|\,|\mathbf{v}_{\mathrm{LG}}|}\right), \tag{4}$$

with uncertainty propagated over posterior density realizations, bias corrections, and the assumed conversion from tracer density to mass density.

These questions produce the following hypotheses:

1. H_1: the Great Attractor wall has a posterior-supported continuation through the inner ZoA after the selection function is modeled.
2. H_2: the Vela data favor a two-wall or branched topology over a single smooth wall.
3. H_3: ZoA completion changes basin boundaries most strongly near the plane and near the inferred watershed, while leaving robust high-contrast structures stable.
4. H_4: including the ZoA constraints changes the amplitude and/or direction of the reconstructed Local Group acceleration by more than the uncertainty induced by catalogue perturbations alone.

The hypotheses are intentionally allowed to fail. A null result would demonstrate that current basin or acceleration estimates are robust to the available ZoA information, which is itself an important cosmographic result.

4 Data architecture

4.1 Master-catalogue design

The first deliverable is a master catalogue in Galactic coordinates. Every record will retain an immutable source identifier and a provenance record so that duplicate detections, cross-matches, and conflicting measurements can be audited.

Table 1: Initial data layers proposed for Z–ZOA.

Layer	**Primary role**	**Principal limitations to preserve**
HIZOA / Parkes	Southern blind H I baseline; $cz \lesssim 12{,}000$ km s^{-1}	Single-dish beam confusion, flux and linewidth completeness, radio continuum effects, H I population bias.
MeerKAT Great Attractor SMGPS	Inner-plane high-resolution H I constraints	Narrow latitude strip, spatially varying sensitivity, continuum background, candidate reliability.
MeerKAT Vela–SMGPS and Vela–H I	Vela structures through and adjacent to the plane	Mosaic footprint, velocity-resolution limits, selection toward gas-rich spirals, survey-specific depth.
2MASS / 2MZOAG	Near-infrared positions and bright-galaxy context	Extinction correction, stellar crowding, source blending, incomplete spectroscopy.
WALLABY / ASKAP	Broad southern H I context and future expansion	Evolving footprint, release state, sensitivity variation, internal versus public data access.
Cosmicflows-4	Independent distances, groups, and peculiar velocities	Heterogeneous distance indicators, correlated errors, grouping assumptions, spatially nonuniform sampling.

Each catalogue row should contain at least

$$\{\ell, b, cz, \sigma_{cz}, d, \sigma_d, S_{\mathrm{H\,I}}, W_{50}, m_{K_s}, \mathrm{type}, \mathrm{survey}, \mathrm{quality}\}\,, \tag{5}$$

where missing values are represented as missing observations with an explicit reason, not as zeros. Cross-matched records will retain the original measurements as well as the adopted posterior combination. The H I flux and linewidth fields are retained because they support both detection modeling and later Tully–Fisher or mass-estimation analyses.

4.2 Coordinate, velocity, and distance conventions

All sky positions will be transformed to Galactic longitude and latitude (ℓ, b), while the original equatorial coordinates are retained for reproducibility. The analysis will record whether velocities are heliocentric, Local Group corrected, or expressed in another frame. The initial atlas will use heliocentric optical-convention velocities, matching the convention used in the MeerKAT Vela publication (Rajohnson et al., 2024a); any conversion to the CMB or Local Group frame will be a separate, versioned transformation.

For visualization, redshift-space coordinates will be distinguished from distance-space coordinates. A galaxy with no independent distance will not be assigned an artificial distance uncertainty narrower than its redshift-space information supports. Where an external distance posterior exists, it will be sampled directly rather than replaced by a point estimate.

5 Probabilistic reconstruction framework

5.1 Selection function

Let $C_s(\ell, b, d, M, T)$ denote the probability that a galaxy with location (ℓ, b, d), luminosity or mass proxy M, and latent type T is detected by survey s. The joint completeness is not a simple product when catalogues share confusion, masking, or cross-match requirements, but a first implementation can represent the survey layers through conditional detection models:

$$C_s = P(\text{detected} \mid \ell, b, d, M, T, \mathcal{I}_s), \tag{6}$$

where $\mathcal{I}_s$ includes sensitivity, beam size, velocity resolution, continuum background, stellar density, extinction, and footprint information.

The observed count in a voxel is then modeled as a thinned realization of an underlying population. In a simple Poisson formulation,

$$N_{s,j} \sim \text{Poisson}\left(\int_{V_j} C_s(\mathbf{x}, M, T)\, \lambda(\mathbf{x}, M, T)\, \mathrm{d}V\, \mathrm{d}M\, \mathrm{d}T \right), \tag{7}$$

where λ is the latent intensity and V_j is voxel j. Equation 7 makes the key distinction explicit: a low observed count can arise from low density, low completeness, or both.

Selection functions will be estimated or bounded using published completeness limits, injection/recovery tests where available, overlap comparisons, and control regions outside the most severe obscuration. The paper's first release should publish the assumed selection-function parameters and sensitivity analyses rather than a single opaque completeness correction.

5.2 Latent density field

The latent tracer density $\delta_g(\mathbf{x})$ will be represented on an adaptive three-dimensional grid or a set of basis functions. The posterior can be written schematically as

$$P(\delta_g, \boldsymbol{\eta} \mid \mathcal{D}) \propto P(\mathcal{D} \mid \delta_g, \boldsymbol{\eta}) P(\delta_g \mid \boldsymbol{\eta}) P(\boldsymbol{\eta}), \tag{8}$$

where $\mathcal{D}$ is the combined catalogue, $\boldsymbol{\eta}$ contains selection, bias, and nuisance parameters, and the likelihood includes detections, upper limits, and uncertainty distributions.

The initial implementation should use a conservative prior that allows walls, filaments, knots, and voids without forcing a particular topology. Gaussian-process or lognormal-field priors are

possible starting points, but the choice must be tested against mock catalogues passed through the same survey masks. The posterior should produce an ensemble of density realizations, not a single smoothed map.

5.3 Cross-survey association

Cross-matching will be treated as an inference problem. Candidate associations will use angular separation, velocity consistency, H I parameter agreement, infrared morphology, and survey-specific positional uncertainties. A match probability p_{ij} will be retained for each candidate pair. Conflicting matches will not be silently resolved by choosing the nearest source.

The resulting catalogue will distinguish three states:

1. a confirmed cross-survey association;
2. a probabilistic association with multiple plausible counterparts;
3. an unmatched detection that remains scientifically useful.

This is important because the absence of an optical or infrared counterpart is itself informative about obscuration and population selection, but it is not proof that the radio detection is spurious.

5.4 From density to velocity

For each posterior density realization, a velocity field can be inferred using a chosen bias and dynamical prescription. In linear theory, the peculiar velocity is related schematically to the density contrast by

$$\mathbf{v}(\mathbf{x}) \propto H_0 f(\Omega_m) \int \delta_m(\mathbf{x}') \frac{\mathbf{x}' - \mathbf{x}}{|\mathbf{x}' - \mathbf{x}|^3} \, \mathrm{d}^3\mathbf{x}', \tag{9}$$

up to normalization, boundary conditions, smoothing, and the relation between galaxy and matter density. The implementation will compare at least two reconstruction prescriptions or smoothing scales to ensure that topological claims are not artifacts of one dynamical approximation.

Measured peculiar velocities from Cosmicflows-4 will enter as noisy constraints on the line-of-sight component. The model will account for distance-modulus uncertainty, group correlations, and the fact that the CF4 sample is not spatially uniform. Redshift-only galaxies will still constrain the density field, but they will not be treated as direct peculiar-velocity measurements.

5.5 Topology and watershed comparison

The atlas will identify walls and filaments using a reproducible ridge or skeletonization algorithm applied to each posterior realization. A candidate structure will be summarized by its posterior occupancy, ridge persistence, velocity interval, and connectivity across the plane.

For basin analysis, streamlines of the reconstructed velocity field will be integrated from a common set of seed points. Basin labels will be assigned by the terminal attractor or sink under a fixed integration rule. The baseline and ZoA-informed runs will be compared using:

- the displacement of basin boundaries in physical units;
- the fraction of voxels that change basin assignment;
- a label-invariant Jaccard overlap for matched basin regions;

- the posterior probability that a streamline crosses a specified boundary;
- the dependence of all changes on Galactic latitude, distance, and density contrast.

Because watersheds are sensitive to smoothing and boundary conditions, every headline topology result will be accompanied by a robustness panel over reconstruction choices.

5.6 Local Group acceleration

The reconstructed acceleration will be evaluated from the same posterior density ensemble used for the flow field. A dipole-like estimator can be written schematically as

$$\mathbf{g}_{\rm rec}(R) = A \int_{|\mathbf{x}|<R} \delta_m(\mathbf{x}) \frac{\mathbf{x}}{|\mathbf{x}|^3}\, \mathrm{d}^3\mathbf{x}, \tag{10}$$

where A absorbs the cosmological normalization and bias mapping. The analysis will report the direction, amplitude, convergence with radius, and change after adding each data layer. Equation 4 will quantify the directional agreement with the adopted Local Group velocity vector.

The result will not be interpreted as a unique mass measurement. H I selection, galaxy bias, unobserved low-surface-brightness populations, external mass beyond the analysis volume, and nonlinear dynamics all contribute uncertainty. The defensible question is whether the ZoA data shift the posterior in a way that is larger than those known uncertainties.

6 Staged analysis plan

6.1 Run 1: nearby structures

The first run will use

$$0 < cz < 12{,}000 \ \mathrm{km\,s^{-1}}, \tag{11}$$

anchored by HIZOA and supplemented by available MeerKAT, 2MASS, optical, and Cosmicflows-4 data. This volume contains the Local Void, Hydra–Centaurus structures, the Norma/Great Attractor region, and nearby walls. It is small enough to validate the catalogue and selection-function machinery before deeper reconstruction.

6.2 Run 2: Vela and deeper Great Attractor structure

The second run will extend to

$$0 < cz < 25{,}000 \ \mathrm{km\,s^{-1}}, \tag{12}$$

to include the deeper MeerKAT Great Attractor and Vela structures. The Vela–H I publication explicitly combines its survey with Vela–SMGPS and adjacent optical/NIR data to trace structures to approximately 25,000 $\mathrm{km\,s^{-1}}$ (Rajohnson et al., 2024a). This stage is the principal test of whether multiple walls and possible plane crossings are part of a connected three-dimensional web.

6.3 Run 3: extended cosmographic context

Only after Runs 1 and 2 pass validation should the analysis be extended to the full usable range of the chosen Cosmicflows and H I data. The wider run will test whether ZoA-induced changes remain local or propagate into the inferred convergence of the bulk flow. A larger volume must not be treated as automatically better: it also introduces more incomplete survey layers, stronger selection variation, and greater sensitivity to external boundary conditions.

7 Validation and falsification

The project will be evaluated with tests designed to expose overinterpretation.

7.1 Mock-catalogue recovery

Synthetic density fields containing known walls, filaments, and voids will be passed through the published or assumed survey masks. The reconstruction must recover the injected structures with calibrated coverage. A method that produces visually plausible structures in a masked mock but fails to recover known topology will not be used for the science claims.

7.2 Leave-one-survey-out tests

The atlas will be reconstructed repeatedly after removing HIZOA, MeerKAT, 2MASS, or the relevant WALLABY layer. The changes attributable to each survey will be reported separately. This guards against the appearance that one survey confirms a structure that was effectively imposed by another.

7.3 Negative controls

Control regions at comparable redshift and survey sensitivity but outside the most severe ZoA will be analyzed with the same pipeline. If the method finds an excess of wall connectivity or basin instability in every control region, the result is likely methodological rather than astrophysical.

7.4 Catalogue perturbations

Positions, velocities, linewidths, distances, and match probabilities will be perturbed according to their reported uncertainties. The posterior distribution of topology and acceleration statistics will be compared with the shift caused by adding the ZoA data. This is the direct test required by hypotheses H_3 and H_4.

8 A scientific visualization, not a substitute for inference

The public-facing atlas should make the observational logic visible without implying that the visualization itself is the measurement. A useful interface would provide synchronized layers:

1. observed galaxies in redshift space;
2. H I detections, with survey provenance and detection uncertainty;
3. near-infrared and optical counterparts;
4. peculiar-velocity constraints;
5. the survey-completeness field;
6. posterior density realizations in the masked region;
7. velocity streamlines and watershed basins.

The Milky Way plane can be shown as a translucent foreground obstruction, but the interface should also display the selection function so that a dark band is not visually confused with a void. A user should be able to toggle from a point estimate to an ensemble of posterior maps and inspect how a claimed wall or basin boundary changes across realizations.

The visualization is therefore a diagnostic and communication layer. It must never be used to select a preferred reconstruction after viewing it. All structure-detection thresholds, smoothing scales, and basin rules should be fixed in the analysis configuration before the final visualization is rendered.

9 Expected contribution

If completed, Z–ZOA would contribute four kinds of result.

First, it would provide a harmonized catalogue in which H I, near-infrared, optical, and distance data can be queried with provenance and uncertainty preserved. Second, it would provide a quantitative selection-function volume that identifies where the map is constrained and where it remains weakly informed. Third, it would turn the Great Attractor and Vela questions into posterior topology tests rather than isolated catalogue overlays. Fourth, it would measure the dynamical consequence of ZoA completion for basin boundaries and Local Group acceleration.

The contribution is intentionally complementary to the MeerKAT surveys. The surveys provide new observations and structure-specific analyses; Z–ZOA would provide a cross-survey inference layer that asks how those observations change a common three-dimensional and dynamical model. The framework also makes room for future releases without requiring the analysis to be redesigned from scratch.

10 Limitations and responsible interpretation

Several limitations should be explicit.

1. H I selection is not a census of all galaxies. The reconstructed density field is a model of the underlying population conditioned on a gas-rich tracer and complementary catalogues.

2. Redshift space is not real space. Fingers of God, coherent infall, and peculiar velocities require a model; they cannot be removed by a purely visual coordinate transform.

3. Distance measurements are heterogeneous. Cosmicflows-4 combines multiple distance indicators and groups; correlations and systematic uncertainties must be retained.

4. A named supercluster is not necessarily a gravitationally bound object. The paper uses walls, filaments, overdensities, flow basins, and watershed boundaries as operational definitions.

5. A better map does not guarantee a unique explanation of the Local Group motion. External mass, bias, nonlinear dynamics, and the adopted reference frame remain relevant.

6. Current WALLABY numbers and access conditions are time-dependent. The analysis must record the exact release, footprint, and access state used for any numerical result.

The manuscript therefore avoids the claim that nobody has mapped the ZoA. Its claim is narrower and testable: modern observations now make it possible to unify several partial views and quantify how their incompleteness affects three-dimensional structure and dynamics.

11 Reproducibility and data-release plan

The project should be released in versioned stages. The first public software release should include catalogue-ingestion schemas, coordinate and velocity-convention metadata, cross-match diagnostics, survey masks, selection-function parameter files, and mock-catalogue generators. The second should include the density and velocity reconstruction configuration. The final release should include posterior summaries, basin-comparison tables, and the exact visualization state used for figures.

Where collaboration agreements limit direct redistribution, the project should publish stable identifiers, download instructions, checksums, and transformation code rather than copying restricted data. Every derived table should record the parent catalogue version and software commit. This is particularly important for evolving survey products such as WALLABY.

12 Conclusion

The Zone of Avoidance is best understood as a structured uncertainty problem. The relevant question is not whether galaxies exist behind the Milky Way; dedicated H I and near-infrared work has already demonstrated that they do. The relevant question is how much the remaining observational incompleteness changes our reconstruction of the local cosmic web and its velocity field.

Z–ZOA proposes a focused answer. It will combine HIZOA, MeerKAT Great Attractor and Vela observations, 2MASS-based catalogues, WALLABY context, and Cosmicflows-4 distances in a master catalogue with explicit provenance and selection functions. It will then propagate uncertainty into density, topology, velocity basins, and Local Group acceleration. The first scientific products should be the Run 1 and Run 2 reconstructions, followed by an evidence-based assessment of Great Attractor continuity, Vela morphology, basin sensitivity, and residual gravitational influence.

The project is deliberately ambitious in its scientific question but conservative in its claims. It does not promise that the ZoA will disappear, nor that one reconstruction will settle the definition of Laniakea or Shapley. It aims instead to make the hidden part of the map measurable, auditable, and dynamically testable.

Acknowledgements

This concept was developed by Zhane Grey as part of the Z–OA project. The proposed work is built on publicly described survey products and published research by the HIZOA, MeerKAT, 2MASS, Cosmicflows, and WALLABY collaborations. The author thanks those teams for making the underlying observations, catalogues, and documentation available.

References

Jonas, J. L. & MeerKAT Team 2016, in *MeerKAT Science: On the Pathway to the SKA*, ed. R. Taylor et al., 1, 1.

Rajohnson, S. H. A., Kraan-Korteweg, R. C., Frank, B. S., et al. 2024, *MNRAS*, 535, 3429, doi:10.1093/mnras/stae2552.

Rajohnson, S. H. A., Kraan-Korteweg, R. C., Frank, B. S., Staveley-Smith, L., & Pisano, D. J. 2026, in *Proceedings of the International Astronomical Union*, 20, 126, doi:10.1017/S1743921325001309.

Schröder, A. C., van Driel, W., & Kraan-Korteweg, R. C. 2019, *MNRAS*, 482, 5167, doi:10.1093/mnras/sty3022.

Staveley-Smith, L., Kraan-Korteweg, R. C., Schröder, A. C., et al. 2016, *AJ*, 151, 52, doi:10.3847/0004-6256/151/3/52.

Steyn, N., Kraan-Korteweg, R. C., Rajohnson, S. H. A., et al. 2024, *MNRAS Letters*, 529, L88, doi:10.1093/mnrasl/slad196.

Tully, R. B., Courtois, H. M., Hoffman, Y., & Pomarède, D. 2014, *Nature*, 513, 71, doi:10.1038/nature13674.

Tully, R. B., Kourkchi, E., Courtois, H. M., et al. 2023, *ApJ*, 944, 94, arXiv:2209.11238.

WALLABY Collaboration 2021, *WALLABY Survey Proposal and Survey Description*, wallaby-survey.org/proposal.

WALLABY Collaboration 2026, *WALLABY Newsletter 24*, February 2026, wallaby-survey.org/newsletters/newsletter-24.