L A B O R A T O R I O S   A L E X A N D R I A

AP-012

Rarity 0.99: The Epistemic Barrier to AI Transfer Between Terrestrial and Planetary Domains

June 2026

Abstract

Transfer learning has become a default assumption in contemporary artificial intelligence: train on domain A, deploy on domain B, adjust as needed. This assumption holds reasonably well across many terrestrial applications. But what happens when domain B is the magnetosphere of Jupiter? This paper presents findings from Alexandria's epistemic deliberation framework on the transferability of AI techniques between terrestrial and planetary science domains. The analysis identifies a "rarity 0.99" barrier—a near-total failure rate for direct, unmodified transfer of AI methods from Earth-based applications to planetary physics. The barrier is not technical but structural: it arises from divergences in the foundational assumptions that underpin each domain, including continuity versus discretization, determinism versus irreducible uncertainty, and bounded versus unbounded parameter spaces. Using NASA's Juno mission as a concrete case study, we demonstrate that computer vision models trained on terrestrial data fail categorically when applied to magnetospheric field reconstruction—not because the models are inadequate, but because the assumptions they encode are inapplicable. We argue that this structural incompatibility represents a broader and largely unexamined phenomenon in cross-domain AI deployment.

Keywords: transfer learning, epistemic barriers, planetary AI, magnetospheric reconstruction, domain divergence, Juno mission

© 2026 Laboratorios Alexandria. All rights reserved. Methodologies described are proprietary trade secrets.

1. Introduction

The promise of transfer learning rests on a deceptively simple premise: knowledge acquired in one domain can be productively applied to another. In natural language processing, models trained on web text perform well on medical records. In computer vision, architectures learned on ImageNet transfer to satellite imagery, histopathology, and industrial inspection. The implicit assumption is that the underlying statistical regularities are sufficiently shared across domains to make transfer not just possible but efficient.

This paper challenges that assumption at its boundary. Specifically, we examine the transfer of AI techniques from terrestrial applications—where models are developed, trained, and validated—to planetary science applications, where the physical phenomena under study operate under fundamentally different regimes. Our analysis emerges from Alexandria's epistemic deliberation framework, which systematically evaluates cross-domain knowledge transfer by identifying structural convergences and divergences between domains.

The central finding is stark. Through the ALETHEIA validation pipeline, which assigns confidence grades to cross-domain correlations via multi-agent deliberation, we identify a "rarity 0.99" in the direct transferability of AI between terrestrial and planetary domains. This metric does not describe a probability in the frequentist sense. It is an epistemic rarity score: among all plausible cross-domain transfers evaluated, the fraction that succeed without deep structural adaptation approaches zero.

The case study that anchors this analysis is NASA's Juno mission to Jupiter. Juno's magnetometer suite produces data on Jupiter's magnetic field that superficially resembles the kind of spatial data that computer vision systems routinely process on Earth. The temptation to apply terrestrial computer vision pipelines—trained on images with continuous spatial gradients, stable lighting, and Euclidean geometry—to magnetospheric reconstruction is understandable. It is also, we argue, structurally misguided.

2. The Assumption Divergence Framework

To understand why transfer fails, it is necessary to look beneath the surface similarity of the data and examine the assumptions that each domain encodes. We identify three primary axes of divergence between terrestrial AI applications and planetary physics.

2.1 Continuity vs. Discretization

Terrestrial computer vision operates on the assumption of spatial continuity. Images are dense pixel grids where adjacent values are correlated. Convolutional architectures exploit this continuity through local receptive fields and translational equivariance. Planetary magnetospheric data, by contrast, is fundamentally sparse. Juno's magnetometer samples Jupiter's field along a single orbital trajectory, producing one-dimensional slices through a three-dimensional, time-varying field. The data is not merely incomplete—it is structurally incompatible with the continuity assumptions that convolutional architectures require.

2.2 Determinism vs. Irreducible Uncertainty

© 2026 Laboratorios Alexandria. All rights reserved. Methodologies described are proprietary trade secrets.

Many terrestrial AI applications operate in deterministic or quasi-deterministic regimes. Object detection assumes that the objects being detected have stable, repeatable visual signatures. Medical imaging assumes that pathological features manifest consistently across patients. Jupiter's magnetosphere is a magnetohydrodynamic system driven by the planet's rapid rotation, internal dynamo, and interaction with the solar wind. The system exhibits chaotic behavior at multiple scales. The "ground truth" is not a fixed label but a distribution of possible states, many of which have never been observed.

2.3 Bounded vs. Unbounded Parameter Spaces

Terrestrial models typically operate within bounded parameter spaces. Pixel values range from 0 to 255. Temperatures vary within known limits. Distances are constrained by the physical setup. In magnetospheric reconstruction, field strengths span orders of magnitude, spatial scales range from planetary radii to ion gyroradii, and temporal dynamics range from the planet's rotation period (roughly ten hours) to sub-second plasma oscillations. Models trained on bounded terrestrial data lack the representational capacity to handle these extreme dynamic ranges without fundamental architectural modifications.

3. Case Study: Juno and the Magnetospheric Wall

The Juno spacecraft, in orbit around Jupiter since July 2016, carries a dual magnetometer system (MAG) that measures the vector magnetic field along Juno's polar orbit. The science objective is to characterize Jupiter's internal magnetic field and understand the dynamics of its magnetosphere—the largest coherent structure in the solar system aside from the heliosphere itself.

Reconstructing Jupiter's magnetic field from Juno data is a classic inverse problem. Given sparse measurements along orbital tracks, the task is to infer the three-dimensional field structure. On Earth, analogous inverse problems in medical imaging (CT, MRI) and geophysics (seismic tomography) have been successfully addressed with deep learning approaches. The question is whether these approaches transfer.

Our analysis, validated at the Structural Convergence level (ALETHEIA Grade A, confidence 0.92), indicates that they do not—at least not without adaptation so profound that it constitutes a fundamentally new methodology rather than a transfer of an existing one. The divergences identified in Section 2 are not abstract concerns but concrete barriers:

Sampling geometry. Juno's highly elliptical polar orbit provides dense sampling near perijove (closest approach, approximately 4,000 km above Jupiter's cloud tops) but extremely sparse coverage at apojove. This is not analogous to undersampling in medical CT, where the geometry is controlled and the object is approximately stationary. The "object"—Jupiter's magnetosphere—rotates, breathes with the solar wind, and restructures itself on timescales shorter than a single orbit.

Absence of labeled data. Terrestrial computer vision benefits from millions of labeled examples. There is no equivalent for planetary magnetospheres. Each planet's magnetic environment is unique. Jupiter is not Earth with a bigger field; its magnetosphere is qualitatively different, dominated by iogenic plasma from Io's volcanic activity—a feature with no terrestrial analog.

© 2026 Laboratorios Alexandria. All rights reserved. Methodologies described are proprietary trade secrets.

Validation impossibility. On Earth, a model's reconstruction can be validated against independent measurements. For Jupiter's deep interior field, there is no independent ground truth. Validation must rely on self-consistency checks and predictions of future observations—a fundamentally different epistemic regime than the train/test split paradigm that underpins most AI evaluation.

4. Rarity 0.99 as Epistemic Metric

The "rarity 0.99" designation requires clarification. It is not a statistical p-value, nor a percentage derived from controlled experiments. It is an epistemic rarity score assigned through Alexandria's deliberation framework, reflecting the degree to which direct, functionally operative transfer between domains is methodologically unusual.

The score emerges from systematic evaluation of domain pairs. For each potential transfer, the framework assesses: (a) whether the source domain's foundational assumptions hold in the target domain; (b) whether the data geometries are compatible; (c) whether the loss functions and evaluation metrics remain meaningful; and (d) whether the model's inductive biases are constructive rather than destructive in the new context. A rarity of 0.99 indicates that across these dimensions, fewer than 1 in 100 evaluated transfers succeed without deep structural modification.

This framing is deliberately epistemic rather than statistical. A statistical claim would require a well-defined population of transfers and a sampling procedure—neither of which is feasible for cross-domain AI deployment, where each domain pair presents unique challenges. The epistemic framing acknowledges this limitation while still providing a useful signal: direct transfer between domains with divergent foundational assumptions is not merely difficult but structurally rare.

5. Beyond Jupiter: The Generalization Question

If the terrestrial-to-planetary transfer barrier is structural rather than technical, a natural question arises: how many other cross-domain transfers that we routinely perform—or routinely assume will work—are similarly compromised?

Complementary analyses within Alexandria's framework suggest that the phenomenon is not unique to the Space–AI interface. Cross-domain correlations between life sciences and astrophysics, evaluated at high confidence (mean 0.89), reveal analogous barriers in the transfer of biological modeling techniques to astrobiological contexts. Similarly, an analysis of the epistemic limits of computation (confidence 0.95) identifies fundamental constraints on the kinds of knowledge that computational methods can extract from domains whose assumptions diverge from those in which the methods were developed.

The pattern that emerges is consistent: cross-domain AI transfer fails not when the domains are "far apart" in some intuitive sense, but when their foundational assumptions diverge. Two domains may share vocabulary, data formats, and even mathematical formalisms while operating under incompatible assumptions about continuity, determinism, observability, or boundary conditions. The surface similarity masks the structural incompatibility.

This has practical implications. The current practice of evaluating transfer feasibility based on data similarity (do the feature spaces overlap?) or task similarity (are the labels analogous?) may be

© 2026 Laboratorios Alexandria. All rights reserved. Methodologies described are proprietary trade secrets.

systematically insufficient. An assumption-level analysis—examining whether the source domain's foundational premises hold in the target—is a necessary complement.

6. Falsifiability and Open Questions

The central claim of this paper—that direct AI transfer between domains with divergent foundational assumptions is structurally rare—is falsifiable. It would be refuted by a demonstrated case of an AI technique (for example, a computer vision model) being applied without deep architectural adaptation to a planetary science problem (for example, reconstruction of Jupiter's magnetospheric field from Juno data) and achieving performance comparable to purpose-built methods.

Several open questions remain, identified through deliberation within the framework. First, the relative contribution of assumption divergence versus architectural limitation is unresolved. It is possible that the transfer barrier reflects not foundational incompatibility but the current limitations of model architectures, and that future architectures with more flexible inductive biases could overcome it. Second, the operational definition of the rarity score (0.99) is based on structured epistemic analysis rather than controlled experimentation, which introduces ambiguity in contexts that demand strictly quantitative metrics. Third, the generalizability of the finding beyond the space–AI interface—to other domain pairs with divergent assumptions—requires systematic investigation.

These are not weaknesses to be minimized but productive uncertainties. The value of the rarity 0.99 finding lies less in its precision than in the question it forces: if we cannot transfer AI between Earth and Jupiter because the assumptions diverge, where else are we transferring AI between domains whose assumptions diverge—and failing to notice?

7. Conclusions

Transfer learning works when the assumptions of the source domain hold, at least approximately, in the target domain. When they do not—when continuity is replaced by sparsity, determinism by chaos, bounded spaces by extreme dynamic ranges—transfer does not merely degrade. It fails categorically.

The magnetosphere of Jupiter, as observed by NASA's Juno mission, provides a concrete and vivid illustration. Computer vision models that excel on Earth do not merely underperform on magnetospheric reconstruction; they encode assumptions that are structurally inapplicable to the target domain. The rarity 0.99 assigned through the ALETHEIA validation pipeline (Grade A, confidence 0.92) quantifies this: direct, unmodified transfer is not just unlikely but methodologically anomalous.

The deeper implication extends beyond planetary science. Cross-domain AI transfer is routinely evaluated on the basis of surface-level similarity—shared data types, analogous tasks, overlapping vocabularies. This paper argues that a more fundamental evaluation is required: an analysis of whether the assumptions encoded in the model are compatible with the assumptions that govern the target domain. Without this assumption-level scrutiny, we risk not just failed transfers but undetected ones—deployments that appear to work because the evaluation metrics themselves inherit the source domain's assumptions.

© 2026 Laboratorios Alexandria. All rights reserved. Methodologies described are proprietary trade secrets.

The question that should follow every proposed cross-domain AI deployment is not “Is the data similar enough?” but “Are the assumptions compatible?” Until that question becomes standard practice, the epistemic barriers identified here will continue to operate—silently, and at scale.

Laboratorios Alexandria — laboratoriosalexandria.com

© 2026 Laboratorios Alexandria. All rights reserved. Methodologies described are proprietary trade secrets.