

Remo Pareschi
STAKE Lab
University of Molise
86100 Campobasso, Italy

August 11, 2026

The Editors-in-Chief
Neurosymbolic Artificial Intelligence
IOS Press / Sage

Dear Editors,

I am pleased to submit the manuscript “*A Minimal κ - τ Logic for Risk-Sensitive Abduction*” for consideration in *Neurosymbolic Artificial Intelligence*. The manuscript is original, is not under consideration elsewhere, and I welcome the journal’s open and transparent review process.

The paper presents a minimal logical kernel for abductive reasoning under risk, built on two governing parameters: an interaction structure κ over hypotheses, and a normative commitment threshold τ . Its central contribution is a combination that, to the best of my knowledge, existing frameworks do not provide: explicit compositional interaction between retained hypotheses, and an internal, rival-sensitive judgment governing when the reasoning process is entitled to commit. The framework is developed in two modes over a common governance schema—a synthetic mode, in which hypotheses combine into emergent composite explanations, and an analytic mode, in which observed outcomes are decomposed into interacting causal clusters.

In accordance with the journal’s policy on extensions of prior work, the manuscript builds on two identified prior publications:

1. R. Pareschi, *Quantum Abduction: A New Paradigm for Reasoning Under Uncertainty*, Sci 7(4):182, 2025 — which introduced the conceptual paradigm of suspended, interaction-driven abduction, together with an embedding-based implementation;
2. R. Pareschi, *Analytic Abduction: Causal Decomposition and Governed Commitment for Human–AI Coordination*, HAR 2026, Springer LNCS, to appear (preprint: arXiv:2607.14641) — which introduced the analytic mode and its governance regime at the conceptual level.

The present manuscript provides the archival advance the journal requires, going substantially beyond both priors in the following respects:

- a complete typed formal language and semantics: a two-sorted grammar separating score-bearing content from non-nesting threshold judgments, with an explicit separation between fixed abductive models and mutable epistemic states, and a governance policy as a model component;
- a worked-out algebraic theory with full proofs and counterexamples: non-associativity of synthesis, idempotence up to structural equivalence, monotonicity restricted to hereditarily constructive terms, and interaction dominance with exact strictness conditions under clamping;
- graded consequence relations over classes of states sharing both governance thresholds, formalizing suspended derivation as a stable, non-closing inferential status;
- a rival-sensitive commitment condition, strengthening the first-past-threshold rule of the

prior work into a margin-based judgment sensitive to the scores and incompatibilities of active rivals;

- the first formalization of the analytic mode within the same semantic kernel: causal clusters with asserted internal interaction structure, coverage and fit valuations, inter-cluster interaction, and two-level (cluster and factor) governed commitment;
- fully computed worked examples for both modes, running on identical interaction data, including an explicit rival-sensitive commitment check;
- a neurosymbolic positioning that is entirely new relative to both priors: a provenance-aware division of labour in which the epistemic parameters are computationally estimable or supplyable while the normative parameters are institutionally governed, with the boundary drawn at the κ/τ interface.

This last point is also why I consider the journal the natural home for the work. The κ - τ logic is designed as the symbolic half of a neurosymbolic division of labour: its epistemic parameters are quantities that existing implementations already estimate from the geometry of neural sentence embeddings, while its normative parameters are set by governance rather than learned. The paper makes this architectural reading explicit and situates it with respect to current neurosymbolic taxonomies.

This work was supported by Ermete—grant number B39J25000400005—2023 MIMIT (Ministero delle Imprese e del Made in Italy). I declare no conflicts of interest.

Thank you for your consideration.

Sincerely,

A handwritten signature in blue ink, reading "Remo Pareschi", is displayed within a light gray rectangular border.

Remo Pareschi
STAKE Lab, University of Molise
86100 Campobasso, Italy
remo.pareschi@unimol.it