# Agentic Roles for Applied Physical-Intelligence Research: A Field-Study Platform and Orchestrated Routing as a World-Model Grounding Instrument

Gil Raitses

aimez.ai Applied Physical Intelligence

## 1. Motivation and Gap

Ha and Schmidhuber introduced a structure in which a learned world model drives action through three parts, a vision model that compresses observations, a memory model that predicts latent dynamics, and a small controller that selects actions, with the controller trained inside the model's own simulated rollout and the policy transferred back to the environment [5]. Current physical-intelligence systems build on this world-model-to-decision lineage, in which a world model predicts outcomes, a critic judges them, and an actor chooses [6].

That lineage carries an unresolved problem. The work defining the modern recurrent world-model agent states that robustly learning and leveraging world models for strong task performance has been an open problem [6]. The methods that address it assume the model can be improved against a simulator or collected interaction data, a hallucinated dream [5] or a small post-training set of trajectories [1].

Applied physical intelligence in a real city has no such setting. There is no dream to train the decision layer in, and the physical world model is assembled at build time from open data rather than learned against a controllable environment. The decision layer then operates on a model that can diverge from the physical plant, and the size of that divergence bounds decision quality [7]. Build-time assembly inherits the underspecification risk of modern machine learning, where pipelines that match held-out metrics can return predictors that behave very differently in deployment domains, so credibility requires field evaluation rather than nominal success alone [3]. This paper contributes an instrument that grounds a build-time physical world model against field observation and records when the decision layer acts on a model that no longer matches conditions: the field-study platform and a catalog of agentic research roles, with orchestrated routing as the decision surface where divergence becomes visible.

## 2. Background

The world-model-to-decision structure maps onto the pax instrument. Vision, which compresses observations into a state [5], is the physical world model layers assembled from open data: crowd stress, mean radiant temperature from a precomputed SOLWEIG model, and street-tree canopy. Memory, which records and predicts dynamics [5], is the field-study temporal record of sessions, observations, and replays. The controller, which selects actions [5, 6], is the routing decision layer, a Dijkstra solver over five cost encodings.

The decision layer faces two documented constraints. A routing metric computed at build time is correct only while edge costs have not changed since customization ran [2]. When the model diverges from the plant, the decision converges only to a neighborhood of the optimum, with neighborhood size scaling with the mismatch [7]. The thermal layer of the world model is built from a SOLWEIG model that requires contemporaneous meteorological inputs at model-run time, so a build-time composite is itself a source of divergence [9].

The instrument adds an operations layer the world-model literature does not specify. A registry defines a catalog of agentic roles with authority tiers of public, researcher, and keyed: a lab-assist role replays sessions with citations, a dossier-reviewer summarizes audit records, a promotion-orchestrator reports ingestion and snapshot status, and a routing-orchestrator mines intent and selects the next interaction node. Every recorded value carries a provenance label under the entity-activity-agent model [10], distinguishing a retrospective field observation from a prospective-only model estimate [4]. The broader autonomous machine intelligence program motivates why grounding a world model that feeds a decision module matters [8].

## 3. Research Goal and Hypothesis

This project asks whether a field-study platform run by agentic research roles can ground a city-scale physical world model well enough that the decision layer tracks lived conditions, and whether provenance-tracked grounding reduces the decision impact of model-plant divergence relative to a build-time-only model [5, 7]. We hypothesize that field observations reveal a measurable divergence between the assembled world model

layers and lived conditions, that this divergence changes routing decisions on a measurable fraction of cases, and that labeling each value as a field observation or a model estimate makes the divergence auditable rather than silent [4].

## 4. Proposed Experimental Approach

The field-study platform is the instrument and routing is the validation surface. A researcher walks a defined route on the Midtown data surface while the platform records a field session of geolocated observations, timestamps, and replays. The agentic roles then replay each session against the world model layers, attach the model value active at each location and time, compile a decision audit record, and fix the provenance of every compared value before analysis.

Three measurements follow. First, the correspondence between field observations and model-layer values is computed per location and time for crowd stress, mean radiant temperature, and shade. The thermal comparison uses the build-time SOLWEIG composite as the model value and the field reading as the observation, making the meteorological-input divergence explicit [9]. Second, routing is run for the walked origin and destination under the model value and under the field observation, and the fraction of cases whose optimal route changes is recorded as the decision impact of divergence [2, 7]. Third, every compared value is labeled a retrospective field observation or a prospective-only model estimate, and label accuracy is checked against the recorded session [10, 4].

## 5. Expected Outcomes and Significance

If the hypothesis holds, field observations would diverge measurably from the assembled world model layers, largest on the thermal layer where a build-time composite stands in for current conditions [9], and that divergence would change the optimal route on a measurable fraction of walked cases, showing that the decision layer acts on a model that does not match the plant [7]. The provenance labels would let a reviewer separate a recorded observation from a model estimate for every compared value, so the divergence is auditable rather than asserted [4]. If field observations instead match the model within the resolution routing requires, the build-time model is sufficient for this surface and season.

The significance is a working instrument for applied physical intelligence, where a world model must be grounded against the real environment rather than improved inside a simulator [5, 6]. The contribution is threefold: the field-study platform as a grounding apparatus, the catalog of agentic research roles that run it with fixed provenance, and the routing decision layer as the surface where model-plant divergence is measured [2]. This addresses how a world model that feeds a decision module is grounded and kept honest [8], with an auditable open-data instrument rather than a learned simulator [1]. The scope is bounded to one city, one season, and one instrument, with routing as the demonstration domain. This paper sits within a research program on decision-valued maps and representation dependence [11].

## References


[1] Assran, M., Bardes, A., Fan, D., et al., LeCun, Y., Rabbat, M. & Ballas, N. (2025). V-JEPA 2. Self-supervised video models enable understanding, prediction and planning. *arXiv:2506.09985*. https://arxiv.org/abs/2506.09985

[2] Bast, H., Delling, D., Goldberg, A., et al. (2016). Route planning in transportation networks. *arXiv:1504.05140*. https://arxiv.org/abs/1504.05140

[3] D'Amour, A., Heller, K., Moldovan, D., et al., & Sculley, D. (2022). Underspecification presents challenges for credibility in modern machine learning. *Journal of Machine Learning Research, 23*(226), 1 to 61. *arXiv:2011.03395*. https://www.jmlr.org/papers/v23/20-1335.html

[4] Freire, J., Koop, D., Santos, E. & Silva, C. T. (2008). Provenance for computational tasks. A survey. *Computing in Science & Engineering, 10*(3), 11 to 21. https://doi.org/10.1109/MCSE.2008.79

[5] Ha, D. & Schmidhuber, J. (2018). Recurrent world models facilitate policy evolution. *Advances in Neural Information Processing Systems 31*. *arXiv:1803.10122*. https://arxiv.org/abs/1803.10122

[6] Hafner, D., Pasukonis, J., Ba, J. & Lillicrap, T. (2025). Mastering diverse control tasks through world models. *Nature, 640*, 647 to 653. https://doi.org/10.1038/s41586-025-08744-2

[7] Kuntz, S. J. & Rawlings, J. B. (2024). Beyond inherent robustness. Strong stability of model predictive control despite plant-model mismatch. *arXiv:2411.15452*. https://arxiv.org/abs/2411.15452

[8] LeCun, Y. (2022). A path towards autonomous machine intelligence. *OpenReview*. https://openreview.net/forum?id=BZ5a1r-kVsf

[9] Lindberg, F., Holmer, B. & Thorsson, S. (2008). SOLWEIG 1.0. Modelling spatial variations of 3D radiant fluxes and mean radiant temperature in complex urban settings. *International Journal of Biometeorology, 52*(7), 697 to 713. https://doi.org/10.1007/s00484-008-0162-7

[10] Moreau, L. & Missier, P. (2013). PROV-DM. The PROV data model. *W3C Recommendation*. https://www.w3.org/TR/prov-dm/

[11] Raitses, G. (2026). On decision-valued maps and representational dependence. *arXiv:2602.11295*. https://arxiv.org/abs/2602.11295