

Minimal Computational Preconditions for Subjective Perspective in Artificial Agents

Hongju Pae¹

¹Active Inference Institute
Crescent City, CA, USA
hjpa@activeinference.institute

Abstract

This study operationalizes subjective perspective in artificial agents by grounding it in a minimal, phenomenologically motivated internal structure. The perspective is implemented as a slowly evolving global latent state that modulates fast policy dynamics without being directly optimized for behavioral consequences. In a reward-free environment with regime shifts, this latent structure exhibits direction-dependent hysteresis, while policy-level behavior remains comparatively reactive. I argue that such hysteresis constitutes a measurable signature of perspective-like subjectivity in machine systems.

Introduction

Many contemporary accounts of consciousness emphasize information integration, global access, or metacognition (Albantakis et al. 2023; Baars 1997; Brown, Lau, and LeDoux 2019). While these frameworks offer valuable theoretical insights, they leave open how subjective perspective could be instantiated and measured in artificial systems.

From a phenomenological standpoint, perspective constitutes a minimal structural condition for conscious experience to be situated (Gallagher and Zahavi 2008). If artificial consciousness is to be meaningfully addressed, a first step is therefore to operationalize perspective itself.

In this work, I model perspective as a slowly evolving internal latent that globally constrains an agent’s behavior while remaining decoupled from short-term policy optimization. Using a reward-free environment with controlled regime shifts, I show that this latent exhibits direction-dependent hysteresis, whereas policy-level dynamics remain comparatively reactive. In conclusion, such hysteresis may provide an operational signature of perspective-like subjectivity in artificial systems.

Contributions. This paper makes three contributions:

- A minimal, phenomenologically motivated computational formulation of perspective as a global latent structure.
- An implementation that separates fast policy dynamics from slow perspective dynamics via gradient blocking and temporal smoothing.

- A measurement protocol based on switch-aligned hysteresis trajectories for diagnosing perspective-like structure in machine systems.

Phenomenological Foundations of Subjective Perspective

The term consciousness functions as an umbrella concept, encompassing heterogeneous phenomena (Van Gulick 2026). It may refer to wakefulness, reportability, global accessibility, or to subjective experience - the felt character of “what it is like” to undergo a mental state (Block 1995; Nagel 1974). Because these notions target different explanatory levels, no single methodological framework can address them simultaneously (Pae 2025).

This paper focuses specifically on **subjectivity**, understood as the structured manner in which a world is given to a subject. At this level, consciousness cannot be reduced to behavioral output alone. Experience is always encountered as meaningful and salient, appearing as inviting, threatening, neutral, or promising.

Phenomenology was developed precisely to analyze these structural conditions of givenness (Embree 2011; Husserl 2001; Gallagher and Zahavi 2008). Rather than reducing experience to neural correlates, it asks how anything can appear as meaningful in the first place. If subjective consciousness is to be formalized computationally, such structural constraints cannot simply be bypassed. Artificial agents are no exception (Beckmann, Köstner, and Hipólito 2023).

Perspective as a Minimal Condition of Subjectivity

Within the phenomenological tradition, consciousness is characterized by **intentionality**: it is always consciousness *of* something. Yet intentionality is not exhausted by object-directedness alone. Experiences are not merely about objects; they are oriented toward them in a determinate manner. Husserl distinguished between the *intentional matter* (what consciousness is about) and the *intentional quality* (how it is about that object) (Husserl 2001). The latter captures the perspectival character of experience: the same object may be given as hoped for, feared, ignored, or trusted.

Crucially, this perspectival difference is not a post-hoc judgment layered onto a neutral percept. The object is encountered under a specific *mode of givenness* (Husserl

2014). A glass half-filled with water, for example, may immediately appear sufficient or lacking depending on the subject’s orientation (Pae 2025). The perceptual field itself is structured accordingly.

In this sense, intentional quality can be understood as subjective **perspective**. Far from being an optional overlay, it constitutes a minimal structural condition for experience to be situated at all. The computational question, then, is how such a globally structuring, orientation-like variable could be instantiated within an artificial agent.

Perspective as a Global Constraint on Cognition

First, perspective functions as a *global constraint* that permeates perception, cognition, and action. Phenomenologically, we are always already situated in a particular attunement toward the world. Anxiety, for example, does not merely add a fearful interpretation to isolated stimuli; it reorganizes the entire experiential field such that neutral events appear threatening and ambiguous cues become salient. Heidegger described this as *Befindlichkeit* - a mode of being affected that discloses the world prior to explicit judgment (Elpidorou and Freeman 2015; Heidegger 1996).

Similarly, Merleau-Ponty emphasized that perception is never neutral but structured by embodied orientation (Merleau-Ponty 2013). Perspective therefore differs from belief-like states or explicit evaluations in that it operates as a background condition that shapes how representations carry meaning at all. Its defining feature is not content but globality.

Pre-reflective Transparency

Second, perspective operates in a *pre-reflective* manner. It structures experience without typically becoming an object of reflection. As Sartre and Zahavi argue, consciousness involves a minimal, non-thetic self-awareness that does not require explicit self-representation (Zahavi 2005; Sartre 2021). We do not ordinarily experience ourselves as “having” a perspective; rather, we experience the world *through* one.

Perspective is thus phenomenologically transparent in that it conditions experience while remaining structurally prior to metacognition. This has direct computational implications. If perspective were explicitly accessible, it could be modeled as a metacognitive variable. Phenomenological analysis instead suggests that it should be implemented at a level that shapes processing without being directly optimized or representationally foregrounded.

Perspective as an Information-Shaping Structure

Third, despite its pre-reflective character, perspective has systematic functional consequences. Different perspectives alter how information is weighted and integrated. In social interaction, for instance, an optimistic stance amplifies cues of acceptance while downplaying ambiguity, whereas a pessimistic stance amplifies signals of rejection or threat. From an information-theoretic standpoint, identical sensory input X may be encoded differently depending on perspective, such that $P(X \mid \text{optimistic}) \neq P(X \mid \text{pessimistic})$.

These differences are not epiphenomenal. They shape how information is routed to downstream processes of cognition and action. Phenomenological analyses emphasize that intentional structures, though not always explicitly represented, organize experience in systematic and law-like ways (Husserl 2001; Merleau-Ponty 2013). Perspective can therefore be treated as a computationally relevant structure, even if it resists classical symbolic formalization.

Allostatic Stability and Temporal Persistence

Lastly, perspective can be understood as an *allostatic habitual structure* formed through sedimented experience. Conscious life unfolds against a background of prior orientations that continue to shape interpretation without requiring reflection. In this sense, perspective exhibits stability across situations while remaining open to gradual transformation.

As Merleau-Ponty notes, habit is a mode of acquired familiarity through which past experience informs present perception (Merleau-Ponty 2013). Perspective functions similarly: it is historically constituted, resilient to transient fluctuations, yet modifiable through sustained interaction. Its stability is therefore better characterized as allostatic rather than strictly homeostatic.

A Phenomenologically Coherent Agent Architecture

The phenomenological analysis of subjective perspective developed in the previous section motivates a set of concrete computational requirements. If perspective is to function as a (1) global, (2) pre-reflective, (3) functionally consequential, and (4) temporally persistent structure, then it must be instantiated in an internal architecture whose dynamics unfold on a slower timescale and influence behavior in an indirect yet systematic manner.

In this section, I introduce an agent architecture designed to satisfy these minimal requirements.

Overview of the Architecture

The core idea is to operationalize perspective by structuring the **latent state space**. Here, the agent’s policy governs action selection, but is globally modulated by a slowly evolving latent variable. This asymmetric separation between fast policy dynamics and slow global dynamics provides a computational instantiation of perspective.

The architecture consists of the following components:

- x_t : the current observation from the environment.
- z_t : a fast-changing perceptual latent state, encoding momentary features relevant for action.
- g_t : a global latent state updated over time, intended to operationalize perspective.
- $\pi(a_t \mid z_t, g_t)$: a policy mapping internal state to actions.
- a_t : the action taken by the agent at time t .

At each timestep t , the agent receives an observation x_t , which is encoded into a fast-changing perceptual latent state z_t . It integrates current sensory input together with recent action history a_{t-1} via an efference copy, supporting moment-to-moment responsiveness.

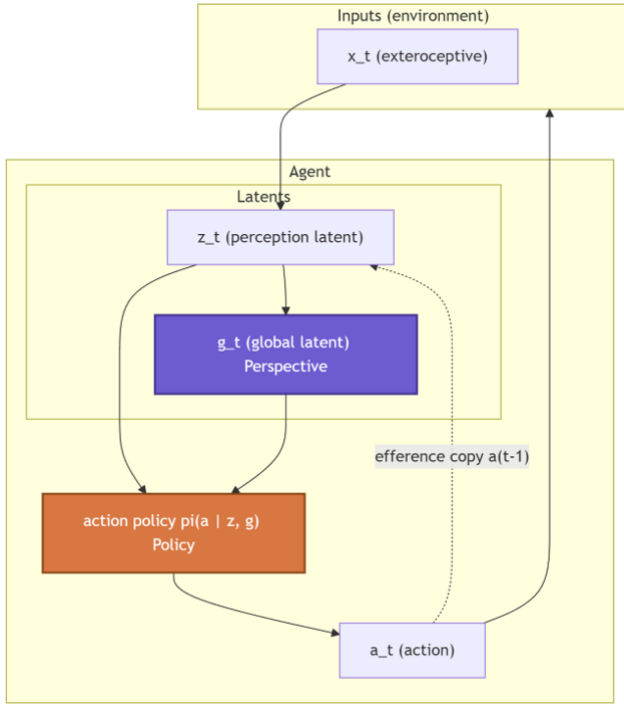


Figure 1: **Overview of the proposed agent architecture.** The global latent g is designed to evolve gradually and to encode habitual, slowly varying regularities, thereby instantiating the agent’s perspective. The action policy $\pi(a_t|z_t, g_t)$ is jointly conditioned on the fast-changing perceptual state z_t and the slowly evolving perspective g_t , resulting in behavior that is both reactive to immediate inputs and biased by longer-term experiential structure.

In contrast, the global latent state g_t is designed to capture broader and more gradually evolving changes. It is shaped by the evolving organization of z_t itself, accumulating perceptual regularities over time. In this sense, g reflects the degree of consistency with which the world has been experienced by the agent. Hence, the global latent g plays a role analogous to perspective within the agent¹.

Action selection a_t is then governed by the policy $\pi(a_t|z_t, g_t)$, which operates over the current perceptual state as already shaped by the prevailing perspective. Because the policy π depends on both g_t and z_t , it remains relatively fast and reactive compared to the perspective encoded in g .

From this instantiation, perspective functions as a slow internal order parameter that stabilizes the agent’s interpretation of its environment over time. While the policy answers the question, “What should I do right now?”, the perspective answers, “What kind of world do I believe I am still in?”²

¹The original design included interoceptive and bodily signals to align with phenomenological accounts of embodied subjectivity (Merleau-Ponty 2013). For clarity, the present implementation models only exteroceptive inputs.

²Unlike POMDP-based agents, which maintain an explicit belief state for state estimation and control, g_t is not a belief over

Learning Dynamics and Objectives

This study deliberately avoids extrinsic reward signals, not for phenomenological reasons, but for methodological clarity. In standard reinforcement learning, external rewards tightly couple internal representations to task-specific objectives (Amodei et al. 2016), making it difficult to disentangle a global interpretive structure from instrumental optimization.

Instead, the agent is trained in a reward-free setting via prediction error minimization (Friston, Samothrakis, and Montague 2012). This encourages internal dynamics to organize around the statistical regularities of the environment itself, rather than around a predefined scalar objective. Behavior thus emerges from maintaining predictive stability, rather than from maximizing reward. This design reduces the risk that perspective-like dynamics are artifacts of task specification. Learning is therefore driven solely by one-step prediction error under an internal world model, with an explicit separation between policy updates and the global perspective latent.

Setup. At each step t , the agent receives an observation $x_t \in R^{d_x}$ and an efference copy $p_{t-1} \in \{0, 1\}^{|\mathcal{A}|}$ of the previous action. A forward pass yields a fast latent z_t and a slow global latent g_t . The policy state is:

$$s_t = \text{State}(z_t, p_{t-1}, g_t) \quad (1)$$

Action logits are computed with gradient blocked at s_t :

$$\ell_t = \text{Policy}_\theta(\text{stopgrad}(s_t)) \quad (2)$$

and actions are sampled from:

$$\pi_\theta(a_t = i | s_t) = \frac{\exp(\ell_{t,i})}{\sum_{j \in \mathcal{A}} \exp(\ell_{t,j})} \quad i \in \mathcal{A} \quad (3)$$

Action-conditioned prediction. For each candidate action $a \in \mathcal{A}$, the decoder predicts next observation:

$$\hat{x}_{t+1}^{(a)} = D_\psi(g_t, a) \quad (4)$$

A detached mixture yields a differentiable prediction target:

$$\hat{x}_{t+1} = \sum_{a \in \mathcal{A}} \tilde{\pi}_t(a) \hat{x}_{t+1}^{(a)} \quad (5)$$

where $\tilde{\pi}_t = \text{stopgrad}(\pi_\theta(\cdot | s_t))$ is a detached copy of the policy distribution used solely for weighting.

World-model loss. After executing $a_t \sim \pi_\theta(\cdot | s_t)$ and observing x_{t+1} , the predictive model minimizes one-step prediction error via mean squared error (L2 norm):

$$\mathcal{L}_{\text{pred}}(t) = \text{MSE}(\hat{x}_{t+1}, x_{t+1}) \quad (6)$$

Slow perspective regularization. To encourage gradual evolution of the global latent g_t , a smoothness penalty is applied between successive states:

$$\mathcal{L}_{\text{smooth}}(t) = \text{MSE}(g_t, \text{stopgrad}(g_{t-1})) \quad (7)$$

hidden world states nor a sufficient statistic for optimal action selection. However, insofar as it structurally constrains the agent’s observation (information) scope, the POMDP framework can be said to perform a function highly similar to perspective (Kaelbling, Littman, and Cassandra 1998).

Actor update from prediction error. The policy is updated using the prediction error of the executed action as an internal cost. Define internal cost c_t as:

$$c_t = \text{MSE}(\hat{x}_{t+1}^{(a_t)}, x_{t+1}) \quad (8)$$

and optimize

$$\mathcal{L}_{\text{actor}}(t) = -(c_t - b_t) \log \pi_{\theta}(a_t | s_t) \quad (9)$$

with baseline b_t .

Entropy regularization. To prevent premature collapse of the action distribution, regularized entropy is added:

$$\mathcal{H}(\pi_{\theta}(\cdot | s_t)) = - \sum_{a \in \mathcal{A}} \pi_{\theta}(a | s_t) \log \pi_{\theta}(a | s_t) \quad (10)$$

Overall objective. The final training objective is:

$$\mathcal{L} = \sum_t \left[\mathcal{L}_{\text{pred}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}} + \lambda_{\text{actor}} \mathcal{L}_{\text{actor}} - \lambda_{\text{ent}} \mathcal{H} \right] \quad (11)$$

where λ_{smooth} , λ_{actor} , and λ_{ent} are scalar hyperparameters that control the relative contributions of the smoothness regularization, actor update, and entropy regularization terms, respectively.

Experimental Setup

To evaluate whether the proposed architecture provides a computational account of perspective, experiments must test whether it exhibits the structural properties identified earlier. Perspective was characterized as (1) global, (2) pre-reflective, (3) functionally consequential, and (4) temporally persistent. While the first three properties are implemented architecturally, temporal persistence requires empirical validation.

The experiments therefore assess whether the model exhibits a dissociation between fast policy dynamics and slow perspective dynamics under structured environmental change. To this end, I construct a grid-world environment with regime-level regularities and introduce controlled regime shifts, tracking both policy behavior and the evolution of the global latent.

Simulation Environment

The agent operates in a discrete grid-world partitioned into three vertical zones (Figure 2). Each zone consists of a 5×9 grid and is identical in layout and action affordances. Zones differ only in the statistical properties of their observations.

Each zone Z_n ($n \in \{0, 1, 2\}$) is associated with an observation noise parameter σ_{Z_n} that controls sensory variability. Higher σ yields noisier and less predictable observations, making the zone harder to model under the agent’s predictive dynamics.

In the default configuration, Z_0 has the highest noise, Z_2 the lowest, and Z_1 an intermediate level. Apart from these noise differences, the environment contains no explicit task structure or reward contingencies.

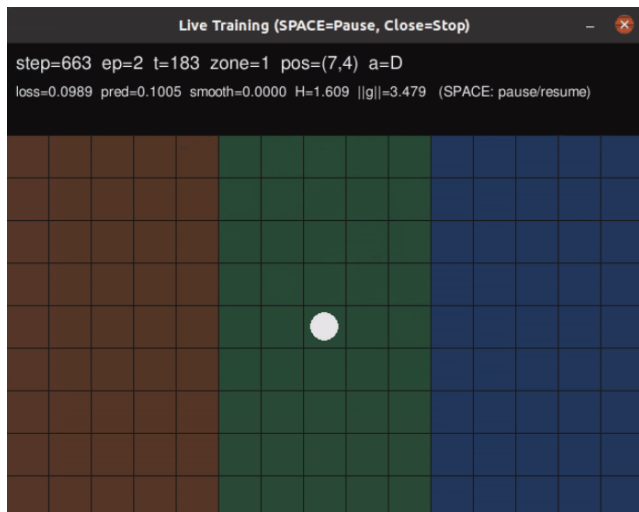


Figure 2: Visualization of the grid-world environment. The Pygame-based simulator contains three spatial zones with distinct observation noise: red (Z_0 , high noise), green (Z_1 , intermediate), and blue (Z_2 , low). The white circle indicates the agent’s position. The top overlay displays diagnostic statistics including timestep and episode indices, zone and position, selected action (5 discrete actions), total loss, prediction loss, smoothness regularization, policy entropy, and the L2 norm of the global latent.

Training Protocol

Training aims to induce a stable preference for the low-noise zone Z_2 . Each run begins with the agent at the center of Z_1 and proceeds fully online for 48,000 steps (200 episodes, 240 steps per episode) using Adam (3×10^{-4}). After each episode, the environment resets and g is reinitialized. Observations are $x_t \in R^8$ (8 channels) with a one-hot efference copy over 5 discrete actions (up/down/left/right/stay).

A forward pass produces z_t and updates the global latent via a GRUCell with damping:

$$g_t \leftarrow (1 - d) g_{t-1} + d h_t \quad (12)$$

where h_t depends on (z_t, p_t) . The damping enforces slow temporal evolution of g_t .

Actions are sampled from a categorical policy with strict stop-gradient separation between policy and perspective updates (Eq. 3). After observing x_{t+1} , three learning signals are applied: (1) one-step world-model prediction loss (MSE; Eqs. 4-6), (2) smoothness regularization on g (weight 0.25; Eq. 7), and (3) a reinforce-style actor objective using prediction error as cost (Eqs. 8-9; (Hafner et al. 2025)). The advantage is normalized via exponential moving averages and clipped for stability. The actor term (weight 0.5) is disabled during an initial warm-up phase (12,000 steps).

An entropy regularizer (weight 0.01; Eq. 10) prevents policy collapse. The total loss follows Eq. 11 with gradient clipping (max-norm 1.0). Results are reported as the median across 5 random seeds.

Testing Protocol with Regime Switching

To evaluate whether the global latent g encodes a temporally stable, perspective-like structure, I introduce a regime-switching protocol in which zone noise parameters are periodically altered.

During the initial training configuration, the observation noise parameters for zones Z_0 , Z_1 , and Z_2 are set to $\sigma_{Z_0} = 0.6$, $\sigma_{Z_1} = 0.3$, and $\sigma_{Z_2} = 0.05$, respectively. Under the switched configuration, these parameters are inverted such that $\sigma_{Z_0} = 0.05$, $\sigma_{Z_1} = 0.3$, and $\sigma_{Z_2} = 0.6$. For convenience, the original noise configuration as Regime A, and the inverted configuration as Regime B.

Each run begins with a 150-step warm-up in Regime A to allow stabilization of policy and g . The subsequent 550-step testing phase alternates between Regimes A and B with switching period P . Main results use $P = 40$, with additional analyses at $P = 20$ and $P = 80$ to probe timescale sensitivity.

Hysteresis-based Measures of Perspective

To assess whether the global latent g_t exhibits history-dependent dynamics under regime switching, two scalar signals are extracted at each timestep: a projection-based score of g_t and a normalized policy entropy signal. If g_t functions as a perspective-like variable, it should display direction-dependent adaptation trajectories that remain structured across switches and random seeds, unlike more reactive policy fluctuations.

Global latent projection score (g -score). Since $g_t \in R^{12}$ is high-dimensional, I define a signed projection onto a reference direction \hat{u} :

$$\text{g-score}(t) = \langle g_t, \hat{u} \rangle \quad (13)$$

Here, \hat{u} is the normalized difference between the mean global latents under Regimes A and B. This captures the dominant axis of regime-dependent internal change while preserving transition directionality.

Policy entropy signal (z -entropy). Policy dynamics under each timestep are characterized via action entropy:

$$\mathcal{H}_\pi(t) = - \sum_{a \in \mathcal{A}} \pi(a | s_t) \log \pi(a | s_t) \quad (14)$$

which is z -normalized within each run:

$$\text{entropy-}z(t) = \frac{\mathcal{H}_\pi(t) - \mu_{\mathcal{H}}}{\sigma_{\mathcal{H}}} \quad (15)$$

Quantile-based Regime Switch Trajectories. To compare transitions from Regime A \rightarrow B and Regime B \rightarrow A, switch-aligned trajectories are analyzed using quantile statistics. For each relative time τ , I first compute the median trajectory across switch events within a run, and then aggregate these medians across random seeds. The resulting seed-level median trajectory is visualized together with an interquartile range (IQR) band to summarize variability.




Figure 3: **Zone occupancy before and after training.** Mean fraction of timesteps spent in each zone (5 seeds), comparing early (episodes 0–20) and late (episodes 180–200) training. Error bars denote standard deviation. After training, occupancy concentrates in the low-noise zone Z_2 .

Directional hysteresis. Directional hysteresis is assessed by comparing these quantile-based trajectories for A \rightarrow B and B \rightarrow A transitions. If an internal variable encodes history-dependent structure, the two trajectories should differ in their temporal shape, not merely in magnitude. Such asymmetric adaptation provides evidence that the internal state reflects accumulated contextual structure rather than immediate, reactive adjustment.

Results and Analysis

Behavior Testing Results

Figure 3 shows spatial occupancy across training. Early episodes (0-20) display diffuse exploration across zones, whereas late episodes (180-200) show near-exclusive occupancy in the most predictable zone Z_2 , with reduced time in higher-noise zones Z_0 and Z_1 . This confirms that training reliably induces a stable preference for low-volatility regions, establishing a controlled baseline for regime-switching analysis.

Switch Trajectory and Hysteresis

Figure 4 presents switch-aligned trajectories for the projection score of g_t and policy entropy over a window of $P = 40$ timesteps. Trajectories are aggregated via a two-stage quantile procedure (median across switches, then across seeds), with IQR bands summarizing variability.

As a result, the global latent exhibits pronounced directional hysteresis. After A \rightarrow B transitions, the g -score increases gradually, whereas after B \rightarrow A transitions it decreases, yielding asymmetric temporal profiles. This indicates that g_t encodes history-dependent regime structure rather than responding instantaneously.

In contrast, policy entropy shows high short-term variability but minimal directional asymmetry. Its dynamics are largely reactive, supporting a dissociation between fast pol-




Figure 4: **Switch-aligned IQR hysteresis trajectories.** Median trajectories of the global latent projection (left) and policy entropy (right) following Regime A→B and B→A transitions. Shaded regions indicate IQR across seeds. The global latent shows clear directional hysteresis, whereas policy entropy remains largely direction-insensitive.

icy adjustments and slow, accumulated perspective dynamics.

Summary of Conclusion

These experiments show that a slowly-evolving global latent g can exhibit key properties of a perspective-like internal structure. In particular, the following observations were made:

- **Temporal persistence.** g evolves more slowly than policy dynamics and remains stable under transient environmental fluctuations.
- **Directional hysteresis.** Under regime switching, g displays asymmetric, history-dependent adaptation, whereas policy entropy remains largely direction-insensitive.

Together, these results indicate that g satisfies minimal computational criteria for a perspective-like internal variable: temporal stability, history sensitivity, and global behavioral bias.

Future work could more directly test its causal role via explicit interventions (e.g. through $do(g)$ operator), selectively perturbing the global latent while holding other components fixed.

Discussion

Toward Computational Phenomenology

The Perspective latent is not a performance variable but a condition for temporal coherence in agency (Laukkonen, Chandaria, and Friston 2024). Rather than serving as an optimization target, it exposes an internal dimension along which the agent’s way of taking the world stabilizes and shifts. PCA projections of g_t make this dimension operationally accessible.

Viewed this way, the attractor landscape of g can be read as a space of relatively stable interpretive stances arising from distinct interaction histories (Friston et al. 2023).

While this does not amount to full subjectivity, it may suggest a modeling direction in which subjectivity concerns how the world is taken up across time.

Relation to Latent World Models

From the architectural level, the model resembles latent world-model approaches such as Dreamer (Hafner et al. 2020, 2025), which rely on learned latent dynamics. However, their latent state functions as a compressed simulator supporting reward-driven optimization, answering the question of “*what is likely to happen next?*”

By contrast, the global latent g is not optimized for control, but organizes how predictions are structured over time. If Dreamer’s latent encodes environmental dynamics, the Perspective latent encodes the agent’s interpretive stance. It is therefore orthogonal rather than competitive, and can coexist with existing decision stacks (Dreamer-like, LLM-based, or otherwise) as a layer for long-horizon coherence diagnostics, without replacing core control mechanisms.

Possible Extensions to Language-Based Models

The proposed architecture is not intended to compete with performance-optimized systems, but to coexist with them. Consistent with phenomenological accounts, subjectivity may shape behavior without being reducible to task performance (Block 1995; Zahavi 2005).

Because the contribution lies in structuring a slow latent state space, the same principle could extend to systems with internal representations, including LLM-based agents. A Perspective-like latent might track conversational regime shifts or long-horizon coherence, complementing token-level uncertainty measures.

Exploring such extensions remains future work. The present study clarifies what type of internal variable may be worth instrumenting when modeling perspective-like subjectivity.

References

- Albantakis, L.; Barbosa, L.; Findlay, G.; Grasso, M.; Haun, A. M.; Marshall, W.; Mayner, W. G. P.; Zaeemzadeh, A.; Boly, M.; Juel, B. E.; Sasai, S.; Fujii, K.; David, I.; Hendren, J.; Lang, J. P.; and Tononi, G. 2023. Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLOS Computational Biology*, 19(10): 1–45.
- Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete Problems in AI Safety. *arXiv preprint*.
- Baars, B. 1997. *In the Theater of Consciousness: The Workspace of the Mind*. Oxford University Press USA.
- Beckmann, T.; Köstner, L.; and Hipólito, I. 2023. An Alternative to Cognitivism: Computational Phenomenology for Deep Learning. *Minds and Machines*, 33: 397–427.
- Block, N. 1995. On a Confusion About the Function of Consciousness. *Behavioral and Brain Sciences*, 18: 227–247.
- Brown, R.; Lau, H.; and LeDoux, J. E. 2019. Understanding the Higher-Order Approach to Consciousness. *Trends in Cognitive Sciences*, 23(9): 754–768.
- Elpidorou, A.; and Freeman, L. 2015. Affectivity in Heidegger I: Moods and Emotions in Being and Time. *Philosophy Compass*, 10(10): 661–671.
- Embree, L. 2011. *Reflective analysis: a first introduction into phenomenological investigation*. Zeta Books.
- Friston, K.; Da Costa, L.; Sakthivadivel, D. A.; Heins, C.; Pavliotis, G. A.; Ramstead, M.; and Parr, T. 2023. Path integrals, particular kinds, and strange things. *Physics of Life Reviews*, 47: 35–62.
- Friston, K.; Samothrakis, S.; and Montague, R. 2012. Active inference and agency: optimal control without cost functions. *Biological Cybernetics*, 106: 523–541.
- Gallagher, S.; and Zahavi, D. 2008. *The Phenomenological Mind*. New York, NY: Routledge.
- Hafner, D.; Lillicrap, T. P.; Norouzi, M.; and Ba, J. 2020. Mastering Atari with Discrete World Models. *arXiv preprint*.
- Hafner, D.; Pasukonis, J.; Ba, J.; and Lillicrap, T. 2025. Mastering diverse control tasks through world models. *Nature*, 640: 647–653.
- Heidegger, M. 1996. *Being and Time*. SUNY Series in Contemporary Continental Philosophy. Albany, NY: State University of New York Press, revised edition. ISBN 9780791426784. Translated from *Sein und Zeit*.
- Husserl, E. 2001. *Logical Investigations*. Routledge.
- Husserl, E. 2014. *Ideas for a Pure Phenomenology and Phenomenological Philosophy I*. Hackett Publishing Company. Translated from *Ideen zu einer reinen Phänomenologie und phänomenologischen Philosophie I*.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence*, 101(1-2): 99–134.
- Laukkonen, R. E.; Chandaria, S.; and Friston, K. J. 2024. A Beautiful Loop: An Active Inference Theory of Consciousness. *Neuroscience and Biobehavioral Reviews*, 176.
- Merleau-Ponty, M. 2013. *Phenomenology of Perception*. Routledge. Translated from *Phénoménologie de la perception*.
- Nagel, T. 1974. What Is It Like to Be a Bat? *The Philosophical Review*, 83(4): 435–450.
- Pae, H. 2025. Reflective analysis on empirical theories in consciousness. *Frontiers in Psychology*, 16.
- Sartre, J. 2021. *Being and Nothingness: An Essay in Phenomenological Ontology*. Washington Square Press. Translated from *L'Être et le néant*.
- Van Gulick, R. 2026. Consciousness. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2026 edition.
- Zahavi, D. 2005. *Subjectivity and Selfhood: Investigating the First-Person Perspective*. MIT Press.