Evolutionary Psychological Structures (EPS): A Neurocomputational Architecture for Value-Grounded, Emotion-Modulated AI

Houston Khanyile
CVEST, South Africa

Correspondence:
Houston Khanyile
CVEST
2152 lowrich Road
Newlands West
Durban
4037
South Africa
Email: houstonmhlongo52@gmail.com

Abstract

Human cognition is shaped by evolved motivational drivers, context-sensitive interpretation processes, and neuromodulator-gated plasticity. Yet current artificial intelligence systems lack biologically grounded mechanisms for computing meaning, prioritizing goals, or generating value-aligned behavior. We introduce *Evolutionary Psychological Structures* (EPS), a neurocomputational architecture that unifies contextual processing, value attribution, motivational drivers, emotional modulation, and long-term cultural learning into a coherent framework grounded in canonical cortical circuitry. EPS integrates (1) Layer 2/3 pyramidal contextual integration, (2) hippocampal–vmPFC value computation, (3) Layer 5 pyramidal driver activation, (4) dopaminergic gain modulation from ventral tegmental area (VTA), and (5) orbitofrontal cortex (OFC)–mediated existential state stabilization, consistent with current systems neuroscience accounts of cortical microcircuits and neuromodulation [Larkum, 2013, Markram et al., 2012, Schultz et al., 1997]. We formalize the *Cultural Matrix* W_{DV}, which maps cognitive values to evolutionary drivers (Belonging, Identity, Power), and introduce a biologically plausible three-factor Hebbian learning rule that binds context, value, and emotional salience [Frémaux and Gerstner, 2014]. We show how EPS can be instantiated within value-grounded machine-intelligence systems, enabling affect-sensitive artificial agents that compute meaning from sensory input. Simulations demonstrate EPS's ability to generate driver-dominance phase diagrams, emotional gain effects, and personality-like drift under long-term learning. We argue that EPS provides a foundation for next-generation AGI architectures capable of culturally adaptive, value-grounded, and psychologically interpretable computation.

Acknowledgements

The author thanks mentors and reviewers for critical feedback that substantially improved the theoretical clarity and biological grounding of this work. This research was conducted independently as part of the broader scientific and technological efforts of CVEST.

Significance

Understanding how humans assign meaning, prioritize goals, and make value-based decisions remains a central challenge across neuroscience, psychology, and artificial intelligence. Current AI systems excel at pattern recognition but lack principled mechanisms for representing motivation, emotion, and culturally shaped values, limiting their interpretability and alignment with human behavior. This work introduces Evolutionary Psychological Structures (EPS), a neurocomputational architecture that formalizes how contextual interpretation, value representation, motivational drivers, and emotion interact to guide behavior. By grounding these processes in established cortical and neuromodulatory mechanisms and expressing them in a mathematically explicit framework, EPS bridges biological cognition and artificial intelligence. The framework generates testable predictions about cultural variation, personality formation, and emotionally modulated learning, while providing a foundation for building AI systems that are more interpretable, value-aware, and human-compatible. EPS therefore advances our understanding of cognition and offers a unifying framework with implications across neuroscience, cognitive science, and machine intelligence.

1 Introduction

Modern artificial intelligence systems excel at pattern recognition, optimization, and prediction, yet fail to reproduce essential dimensions of human cognition: contextual meaning, emotional salience, cultural bias, personality, and motivational structure. Humans do not process the world through raw sensory data alone. Instead, sensory information is transformed through layered neurocomputational operations shaped by millions of years of evolutionary pressures driving belonging, identity formation, and power negotiation [Cosmides and Tooby, 2014, Haidt, 2012, Baumeister and Leary, 1995].

Although existing neuro-inspired models—including predictive coding and active inference [Friston, 2010, Friston et al., 2016], reinforcement learning approaches to reward and decision-making [Dayan and Seymour, 2007, Schultz et al., 1997], and hierarchical control accounts of prefrontal cortex [Badre and D'Esposito, 2010]—capture fragments of this pipeline, they lack a mechanistic understanding of how *values*, *drivers*, and *emotion* interact to produce behavior. AGI research, in turn, typically focuses on scalable function approximation [LeCun, 2022], with limited grounding in the biological mechanisms that determine meaning-making, motivational selection, and personality.

Evolutionary Psychological Structures (EPS) provides a unifying computational architecture that explains how the brain transforms sensory input into psychologically meaningful, culturally biased, and emotionally weighted interpretations that drive behavior. The theory integrates:

- evolutionary motivational drivers (Belonging, Identity, Power),
- contextual interpretation (Layer 2/3 pyramidal computation),
- value attribution (hippocampal–vmPFC interactions),
- driver activation (Layer 5 pyramidal outputs),
- dopaminergic gain modulation (VTA),
- cultural learning via three-factor Hebbian plasticity, and
- existential state stabilization (OFC).

EPS is formalized mathematically and implemented computationally as a layered dynamical system. While EPS can be instantiated within proprietary architectures for value-grounded artificial agents, the present paper focuses on the theoretical, mathematical, and neuroscientific foundations of the framework, together with simulation results that validate its internal dynamics.

2 EPS as a Computational Architecture

EPS decomposes human interpretation into a five-stage transformation pipeline:

1. **Contextual filtering**,
2. **Value mapping**,
3. **Driver activation**,
4. **Emotional modulation**, and
5. **Existential state confirmation**.

These stages form a structured computation:

$$S(t) \to C(t) \to V(t) \to D(t) \to M(t) \to Q(t), \tag{1}$$

where $S(t)$ denotes a stimulus or input representation, $C(t)$ the contextual state, $V(t)$ a value vector, $D(t)$ the driver activations, $M(t)$ a neuromodulatory signal, and $Q(t)$ an existential state estimate. Each stage corresponds to a functionally distinct neural population and is parameterized via mathematically defined dynamical systems.

2.1 Drivers as Evolutionary Attractors

EPS asserts that human motivation is organized around three evolutionarily shaped drivers:

- **Belonging (B):** group safety, cohesion, affiliation.
- **Identity (I):** self-stability, narrative consistency, self-definition.
- **Power (P):** control over resources, outcomes, and constraints imposed by others.

These drivers act as latent attractor states in motivational computation. Different cultures, developmental histories, and environmental contingencies shape the weights with which these drivers are activated in response to particular contextual and value patterns [Schwartz, 2010, Hofstede, 1980].

2.2 Context and Value as Interpretive Layers

Context encodes the interpretive framing extracted by Layer 2/3 pyramidal cells and governed by inhibitory interneurons, consistent with evidence for horizontal integration and local circuit computation in superficial cortical layers [Harris and Mrsic-Flogel, 2013, Larkum, 2013]. Context $C(t)$ binds together sensory, task, and prior knowledge signals into a structured representation.

Value is retrieved through hippocampal–vmPFC interaction, capturing meaning attributes such as safety, fairness, loyalty, threat, and opportunity [Howard et al., 2014, Dolan, 2002]. We denote the resulting value vector as $V(t)$, which can be low-dimensional (e.g., safety vs. threat) or high-dimensional (multiple moral and social dimensions [Haidt, 2012]).

2.3 The Cultural Matrix W_{DV}

A central innovation of EPS is the *Cultural Matrix*:

$$W_{DV} = \begin{bmatrix} w_{B,1} & w_{B,2} & \cdots \\ w_{I,1} & w_{I,2} & \cdots \\ w_{P,1} & w_{P,2} & \cdots \end{bmatrix}, \tag{2}$$

which maps values $\to$ drivers as a learnable transformation. Each column corresponds to a value dimension; each row corresponds to one of the three drivers.

This mapping enables:

- cultural variation in value-to-driver mappings,
- personality differences,
- developmental change,
- sociocultural complexity, and
- explainable biases,

without modifying the underlying biological architecture. Different cultures can therefore be understood as different distributions over W_{DV} matrices, rather than different neuroanatomical blueprints.

3 Biological Grounding

EPS is explicitly grounded in canonical cortical microcircuitry and systems-level loops.

3.0.1 Conflict Resolution Among Situation Classes

Real-world scenes frequently instantiate multiple ancestral situation classes simultaneously (e.g., intra-group competition co-occurring with social exclusion). Therefore, contextual filtering requires an explicit *interpretation-level* conflict resolution mechanism, prior to value mapping and existential state confirmation.

Situation-class hypotheses in L2/3. Let $\mathcal{K} = \{1, \dots, K\}$ index ancestral situation classes (e.g., Predation, Resource Scarcity, Social Exclusion, etc.). We model each class $k \in \mathcal{K}$ as a competing contextual hypothesis encoded by a Layer 2/3 pyramidal population with activity $C_k(t)$.

EEA-constrained situational metadata. Let $\mathbf{m} \in [0,1]^d$ denote the situational metadata vector (e.g., magnitude, immediacy, proximity, agency/control, power differential, predictability, reversibility). For each class k, let $\boldsymbol{\pi}_k \in \mathbb{R}^d$ denote an EEA-grounded projection vector encoding how metadata dimensions modulate existential risk within that class.

Viability-weighted class evidence. Define the class evidence (risk projection) as the inner product

$$I_k \;=\; \langle \mathbf{m}, \boldsymbol{\pi}_k \rangle. \tag{3}$$

To prevent catastrophic misclassification, we introduce a precedence bias $\lambda_k > 0$ reflecting the irreversible fitness cost of misinterpreting class k (e.g., $\lambda_{\text{Predation}}$ large). The effective class input becomes

$$I_k^{\text{eff}} \;=\; \lambda_k \, I_k \;=\; \lambda_k \, \langle \mathbf{m}, \boldsymbol{\pi}_k \rangle. \tag{4}$$

CRF as L2/3 winner-take-all dynamics. Interpretive conflict resolution is implemented by lateral competition in Layer 2/3. A minimal rate model is:

$$\tau_C \frac{dC_k}{dt} \;=\; -C_k + \phi\left(I_k^{\text{eff}} - \sum_{j \neq k} w_{kj}\, C_j + b_k \right), \qquad k \in \mathcal{K}, \tag{5}$$

where $w_{kj} \geq 0$ are lateral inhibitory weights (mediated by interneurons), b_k is a baseline excitability term, and $\phi(\cdot)$ is a rectifying nonlinearity.

Under standard conditions for competitive networks, Eq. (5) converges to a stable attractor in which one hypothesis dominates:

$$k^* \;=\; \arg\max_{k \in \mathcal{K}} \left(I_k^{\text{eff}} \right) \;=\; \arg\max_{k \in \mathcal{K}} \left(\lambda_k \left\langle \mathbf{m}, \boldsymbol{\pi}_k \right\rangle \right), \tag{6}$$

and $C_{k^*} \gg C_j$ for $j \neq k^*$. We refer to Eq. (6) as the *Conflict Resolution Function (CRF).*

Downstream use. The dominant contextual hypothesis k^* (and its stabilized activity C_{k^*}) defines the primary interpretive frame for semantic compression (meaning) and subsequent value mapping. Lower-priority hypotheses are not discarded; they remain latent and may modulate downstream driver dynamics and learning, but only one class governs meaning compression at a given moment.

3.1 Hippocampal CA1/CA3 and vmPFC

Hippocampal circuits and ventromedial prefrontal cortex (vmPFC) jointly construct value representations grounded in learned associations, place and event structure, and inferred significance [Howard et al., 2014]. In EPS, a hippocampal–vmPFC loop computes $V(t)$ as a function of context and long-term memory.

3.2 Layer 5 Pyramidal Neurons

Layer 5 pyramidal neurons transform value-weighted signals into driver outputs and project to subcortical structures and motor systems, consistent with their known role in descending control and integrative output [Badre and D'Esposito, 2010, Doya, 1999]. In EPS, Layer 5 populations compute $D(t)$ from $V(t)$ via W_{DV}.

3.3 GABAergic Interneurons

GABAergic interneurons provide contextual inhibition essential for filtering, gating, and competitive selection in cortical circuits [Markram et al., 2004]. In EPS, interneurons modulate the effective gain and sparsity of $C(t)$, shaping which contextual features are allowed to influence $V(t)$ and $D(t)$.

3.4 VTA Dopaminergic Neurons

VTA dopaminergic neurons provide multiplicative gain modulation and reinforcement learning signals [Schultz, 1998, Dayan and Seymour, 2007]. In EPS, neuromodulation is modeled as:

$$W_{DV}^{\text{eff}}(t) = W_{DV}\big(1 + \alpha_M M(t)\big), \tag{7}$$

where $M(t)$ is an emotional modulation signal and α_M a gain parameter. This formalizes the idea that emotionally salient states transiently amplify or attenuate particular value-to-driver pathways, akin to dopamine-dependent gain control in neural systems.

3.5 Orbitofrontal Cortex (OFC)

Orbitofrontal cortex is implicated in representing abstract outcome value, integrating multi-dimensional signals, and tracking latent states of the environment [Rolls, 2014]. In EPS, OFC-like circuitry is modeled as an existential state module computing $Q(t)$, the agent's current orientation along Survival, Stability, and Quality-of-Life axes.

Overall, this grounding converts EPS from a purely psychological theory into a testable neurocomputational model.

4 Mathematical Framework

EPS is expressed as a set of coupled neural population equations. For simplicity, we model each stage as a rate-based population with first-order dynamics and a nonlinear activation function $\phi(\cdot)$ (e.g., rectified linear or sigmoidal).

4.1 Context Dynamics

Context evolves according to:

$$\tau_C \frac{dC}{dt} = -C + \phi\big(W_{\text{sens}} S(t) + W_{\text{cog}} C(t) + b_C\big), \tag{8}$$

where τ_C is a time constant, W_{sens} encodes feedforward sensory weights, W_{cog} recurrent cognitive weights, and b_C a bias term.

4.2 Value Dynamics

Value representation evolves as:

$$\tau_V \frac{dV}{dt} = -V + \phi\big(W_{VC} C(t) + W_{V\text{rec}} V(t) + b_V\big), \tag{9}$$

where W_{VC} maps context to values, $W_{V\text{rec}}$ captures recurrent value interactions, and b_V is a bias.

4.3 Driver Dynamics

Driver activations are given by:

$$\tau_D \frac{dD}{dt} = -D + \phi\big(W_{DV}^{\text{eff}}(t) V(t)\big), \tag{10}$$

with $D(t) = [D_B(t), D_I(t), D_P(t)]^\top$ and $W_{DV}^{\text{eff}}(t)$ the modulated Cultural Matrix.

4.4 Modulation Dynamics

Emotional modulation is computed as:

$$\tau_M \frac{dM}{dt} = -M + \phi\big(W_{MD}^\top D(t) + b_M\big), \tag{11}$$

with W_{MD} a vector projecting driver activations into a scalar modulation signal $M(t)$, controlling gain in W_{DV}^{eff}.

4.5 Existential State Dynamics

The existential state $Q(t)$ integrates value, driver, and modulation signals:

$$\tau_Q \frac{dQ}{dt} = -Q + \phi\big(W_{QV}V(t) + W_{QD}^{\top}D(t) + W_{QM}M(t)\big), \tag{12}$$

where $Q(t)$ may be a three-dimensional vector representing Survival, Stability, and Quality-of-Life state estimates.

4.6 Three-Factor Hebbian Learning

Long-term adaptation of the Cultural Matrix W_{DV} is governed by a three-factor Hebbian rule [Frémaux and Gerstner, 2014]:

$$\Delta W_{DV}^{ij} = \eta\, M(t)\, C_i(t)\, V_j(t), \tag{13}$$

where η is a learning rate, $C_i(t)$ and $V_j(t)$ are pre-synaptic (context) and post-synaptic (value) components, and $M(t)$ gates plasticity according to emotional salience. This binds emotion, context, and value into long-term cultural and personality memory.

5 Implementation Considerations (High-Level)

The EPS architecture can be implemented in a variety of computational substrates, including rate-based neural simulators, deep learning frameworks, and hybrid symbolic–neural systems. In practice, the system may be discretized and integrated numerically, with each of the dynamical equations implemented as update rules over time steps. The present work focuses on the conceptual architecture and its dynamical properties; specific engineering details of applied instantiations are beyond the scope of this paper and may be realized within proprietary machine-intelligence frameworks.

6 Simulation Methods

To demonstrate that EPS yields coherent and interpretable dynamics, we implemented the above equations as discrete-time rate models using Euler integration with fixed time step Δt. Unless otherwise noted, we set $\tau_C = \tau_V = \tau_D = \tau_M = \tau_Q = 50$ ms and simulated trials of length $T = 2000$ ms.

Inputs $S(t)$ were modeled as simple square pulses or step functions, representing the onset and offset of a salient event. We used low-dimensional value vectors $V(t) \in \mathbb{R}^1$ for illustrative purposes (e.g., a single safety–threat axis), although the framework naturally extends to higher-dimensional value spaces.

We conducted four classes of experiments:

1. single-event pipeline dynamics,
2. driver-dominance phase-space exploration,
3. long-term Cultural Matrix learning, and
4. personality emergence under different environments.

All simulations were implemented in a modular codebase (e.g., in Python with standard numerical libraries). Parameter choices were selected to illustrate qualitative behaviors rather than to fit empirical data.

7 Results

7.1 EPS Pipeline Dynamics (Figure 1)

In the first set of simulations, we examined the time course of the EPS pipeline in response to a transient input stimulus $S(t)$. A square pulse was applied between 200 and 300 ms, with $S(t) = 0$ outside this window and $S(t) = 1$ during it.

Figure 1 illustrates a representative trial. The contextual state $C(t)$ rises rapidly following stimulus onset, reflecting immediate contextual interpretation. The value signal $V(t)$ rises after $C(t)$ with slightly slower dynamics, stabilizing into a value estimate that persists beyond stimulus offset. The driver activations $D_B(t)$, $D_I(t)$, and $D_P(t)$ are then driven by $V(t)$ through W_{DV}, with one or two drivers dominating depending on the Cultural Matrix configuration. The modulation signal $M(t)$ increases monotonically as a function of driver activation, and the existential state $Q(t)$ rises accordingly, reflecting the integrated interpretation of the event.

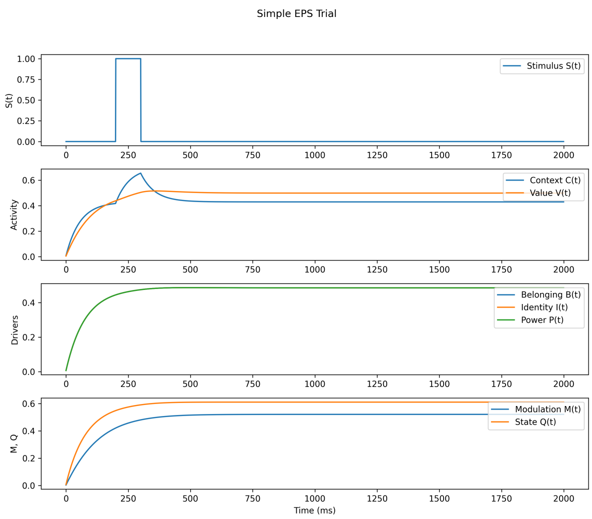

Figure 1: EPS Pipeline Dynamic

7.2 Cultural Effects on Driver Dominance (Figure 2)

We next explored how cultural variation in W_{DV} and neuromodulatory gain α_M interact to shape driver dominance. We parameterized a simple family of Cultural Matrices by varying a Belonging weight parameter along one axis and the gain parameter α_M along another and simulated steady-state driver activations for each pair.

Figure 2 shows a phase diagram of the dominant driver as a function of Belonging weight and modulation gain. For low Belonging weights, Identity tends to dominate; for sufficiently high Belonging weights, Belonging emerges as the dominant driver. Within the parameter ranges tested here, Power did not become the dominant driver, although higher baseline Power weights or stronger gain could produce such regimes. This illustrates how EPS can generate interpretable phase boundaries between motivational regimes based on cultural parameters.

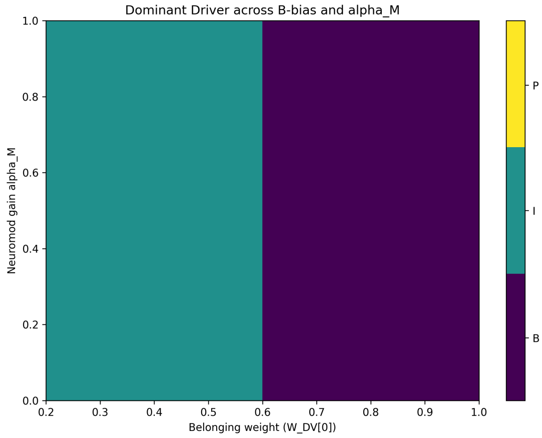

Figure 2: Driver Dominance Phase Diagram

7.3 Cultural Matrix Learning (Figure 3)

To examine long-term Cultural Matrix adaptation, we simulated repeated episodes with the three-factor Hebbian rule updating W_{DV}. Each episode consisted of a brief input, allowing the system to evolve and generate $C(t)$, $V(t)$, $D(t)$, and $M(t)$, after which W_{DV} was updated according to the final or time-averaged states.

Figure 3 plots the trajectories of selected entries of W_{DV} over episodes. With a small learning rate η, the matrix entries drift gradually, reflecting emotionally gated learning: episodes with higher $M(t)$ induce larger updates. This produces slow shifts in value-to-driver mappings, demonstrating how EPS can model cultural learning and the consolidation of motivational tendencies over extended experience.

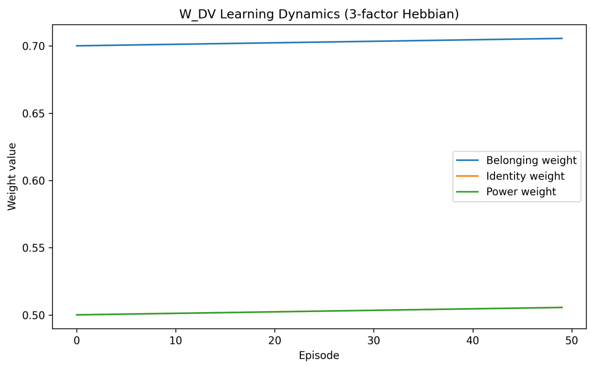

Figure 3: Cultural Matrix Learning Dynamics

7.4 Personality Emergence Under Different Environments (Figure 4)

Finally, we simulated three distinct "environments" characterized by different initial Cultural Matrices and modulation parameters: (1) a cooperative, Belonging-oriented environment, (2) a self-expressive, Identity-oriented environment, and (3) a competitive, Power-oriented environment. For each environment, we ran many episodes and updated W_{DV} via the three-factor rule.

Figure 4 shows the trajectories of the Belonging, Identity, and Power weights in W_{DV} across episodes for each environment. In the cooperative environment, Belonging-related entries strengthen over time; in the self-expressive environment, Identity-related entries become dominant; and in the competitive environment, Power-related entries are increasingly favored. These trajectories illustrate how stable personality-like motivational profiles can emerge from repeated, emotionally salient experience in different environments, without altering the underlying architecture.

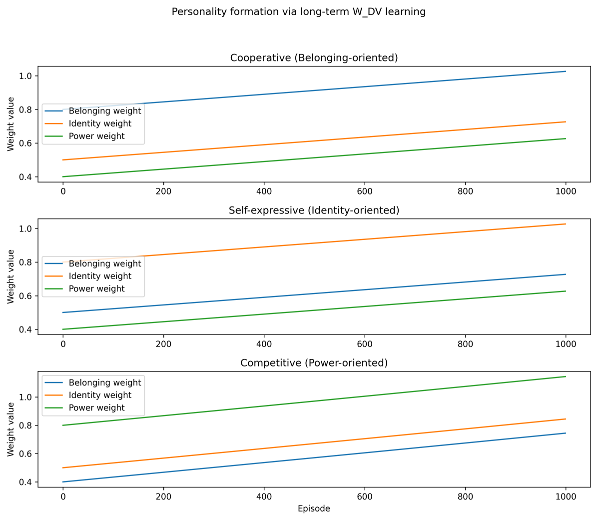

Figure 4: Personality Emergence via Long-Term Learning

Code Availability. All simulation code and reference implementations of the EPS architecture are available at: https://github.com/{{repo-owner}}/eps-neurocomputational-framework

8 Benchmarks and Theoretical Predictions

EPS makes several testable theoretical predictions:

1. **Dopaminergic gain selectively amplifies cultural biases.** Increasing neuromodulatory gain should expand the parameter regions in which particular drivers dominate, in line with known roles of dopamine in gain modulation and salience [Schultz, 1998].

2. **Personality reflects long-run equilibria of Hebbian updating.** Stable individual differences in motivational style (e.g., Belonging-oriented vs. Power-oriented) can be modeled as attractor states of the Cultural Matrix under the three-factor learning rule.

3. **Cultures differ via W_{DV} distributions.** Cross-cultural differences in value-to-driver mappings (e.g., collectivist vs. individualist societies) can be captured as differing priors over W_{DV}, consistent with cultural value orientation work [Schwartz, 2010, Hofstede, 1980].

4. **Emotional salience accelerates learning.** Episodes associated with high $M(t)$ should produce disproportionately large updates to W_{DV}, providing a mechanistic account of emotionally weighted memory and learning.

5. **Value-grounded AI systems will exhibit interpretable motivational structure.** Artificial agents that implement EPS-like architectures should display recognizable motivational profiles and allow post-hoc interpretation of decisions via inspection of W_{DV} and driver trajectories.

These predictions are amenable to empirical testing in both neuroscience and computational cognitive modeling, and to behavioral evaluation in artificial agents.

9 Implications for Artificial General Intelligence

EPS provides a blueprint for integrating meaning, motivation, and emotion into artificial agents. Unlike purely statistical models [Russell and Norvig, 2010, LeCun, 2022], an EPS-based architecture explicitly represents:

- contextual interpretations $C(t)$,
- value vectors $V(t)$,
- motivational drivers $D(t)$,
- emotional modulation signals $M(t)$, and
- existential states $Q(t)$.

This enables:

- meaning-aware computation,
- cultural interpretability,
- emotional grounding,
- psychological explainability,
- value-aligned decision layers, and
- motivational transparency.

Incorporating EPS as an intermediate layer within large-scale AI systems may help bridge the gap between low-level function approximation and high-level human-like reasoning, supporting more aligned, accountable, and controllable AGI.

10 Discussion

EPS unifies evolutionary psychology, cortical computation, neuromodulatory control, and long-term cultural learning into a single mathematical and computational framework. By explicitly modeling drivers, values, context, and emotion, it offers a principled mechanism for meaning, cultural variation, and identity.

In contrast to black-box deep learning systems, EPS emphasizes structured intermediate representations that can be inspected and interpreted. In contrast to purely normative frameworks such as active inference [Friston, 2010], EPS foregrounds evolutionary drivers and culturally learned mappings that shape motivational priorities. These features make EPS a promising foundation for affective, human-aligned AI.

Future work should connect EPS more tightly to detailed biological data, including laminar recordings, neuromodulatory manipulations, and longitudinal studies of personality and cultural learning. On the AI side, integrating EPS with large language models, reinforcement learning agents, and multi-agent environments presents a rich space for experimentation.

11 Limitations

The present formulation has several limitations:

- **Simplified rate models.** We use rate-based dynamics for clarity; spiking models [Buzsáki, 2006, Brunel, 2000] would allow more realistic temporal structure.
- **Low-dimensional value space.** For simplicity, we illustrated simulations with one or a few value dimensions. Real-world values are high-dimensional and context-dependent [Haidt, 2012].
- **Heuristic parameter choices.** The parameters used are heuristic and chosen for qualitative behavior; empirical calibration against neural and behavioral data is an important next step.
- **Abstract existential states.** The Survival–Stability–Quality-of-Life decomposition is a modeling choice; alternative decompositions may better capture human phenomenology.
- **Implementation details omitted.** While EPS can be instantiated within practical AI systems, we do not disclose specific implementation details here, in order to focus on the theoretical and neuroscientific contributions.

These limitations point toward a broader research program rather than weakening the core claims of EPS as a unifying architecture.

12 Conclusion

EPS offers a unified explanation of human motivational computation and provides a computational foundation for emotionally grounded artificial intelligence. Through contextual processing, value attribution, driver activation, neuromodulated gain, and Hebbian cultural learning, EPS captures the mechanisms through which meaning and personality emerge over time.

By formalizing the Cultural Matrix W_{DV} and an emotion-gated learning rule, EPS connects evolutionary drivers, cultural variation, and individual personality into a single framework. We argue that incorporating EPS-like architectures into AI systems is a necessary step toward building value-grounded, culturally adaptive, and psychologically interpretable artificial agents.

Appendix A: Conceptual Overview of Sensia Quotia Computation (SQC)

While the present manuscript focuses on the theoretical and neuroscientific foundations of EPS, it is useful to briefly indicate how EPS can be embedded in applied machine-intelligence systems.

Sensia Quotia Computation (SQC) denotes a family of computational architectures that implement the EPS pipeline as an intermediate layer within artificial agents. At a high level, SQC systems:

- encode context $C(t)$ from sensory or symbolic inputs,
- compute value vectors $V(t)$ using learned semantic and associative representations,
- map values to drivers via a learnable Cultural Matrix W_{DV},
- apply emotional modulation $M(t)$ as gain control over W_{DV}, and
- feed existential states $Q(t)$ into downstream policy or decision modules.

SQC thus operationalizes EPS within agent architectures, allowing agents to exhibit stable motivational profiles, culturally shaped behavior, and emotionally coherent decision-making. Detailed engineering designs, parameterizations, and implementation choices for SQC are proprietary and are not described in this paper; the present appendix is intended only to clarify the conceptual relationship between EPS as a scientific theory and its potential instantiations in artificial intelligence systems.

References

Matthew Larkum. A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151, 2013.

Henry Markram et al. Spike-timing-dependent plasticity: a comprehensive overview. *Frontiers in Synaptic Neuroscience*, 4:2, 2012.

Wolfram Schultz, Peter Dayan, and P. Read Montague. A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599, 1997.

Nicolas Frémaux and Wulfram Gerstner. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in Neural Circuits*, 7:119, 2014.

Leda Cosmides and John Tooby. Evolutionary psychology and the brain. In *The Cambridge Handbook of Cognitive Science*. Cambridge University Press, 2014.

Jonathan Haidt. *The Righteous Mind*. Pantheon, 2012.

Roy F. Baumeister and Mark R. Leary. The need to belong: desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, 117(3):497–529, 1995.

Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

Karl Friston et al. Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68: 862–879, 2016.

Peter Dayan and Ben Seymour. Values and actions in aversion. *Science*, 315(5811):1603–1606, 2007.

David Badre and Mark D'Esposito. Is the rostro-caudal axis of the frontal lobe hierarchical? *Nature Reviews Neuroscience*, 11:719–728, 2010.

Yann LeCun. A path towards autonomous machine intelligence. *Open Review Manuscript*, 2022. arXiv preprint.

Shalom H. Schwartz. Basic values: How they motivate and inhibit prosocial behavior. In *Prosocial Motives, Emotions, and Behavior*. American Psychological Association, 2010.

Geert Hofstede. *Culture's Consequences: International Differences in Work-Related Values*. Sage, 1980.

Kenneth D. Harris and Thomas D. Mrsic-Flogel. Cortical connectivity and sensory coding. *Nature*, 503:51–58, 2013.

Marc W. Howard et al. A unified mathematical framework for hippocampal and neocortical coding of space. *Proceedings of the National Academy of Sciences*, 111(29):E2967–E2976, 2014.

Raymond J. Dolan. Emotion, cognition, and behavior. *Science*, 298(5596):1191–1194, 2002.

Kenji Doya. What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex? *Neural Networks*, 12(7–8):961–974, 1999.

Henry Markram, Maria Toledo-Rodriguez, Yun Wang, Anirudh Gupta, Gilad Silberberg, and Caizhi Wu. Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, 5:793–807, 2004.

Wolfram Schultz. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80:1–27, 1998.

Edmund T. Rolls. *Emotion and Decision-Making Explained*. Oxford University Press, 2014.

Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3 edition, 2010.

György Buzsáki. *Rhythms of the Brain*. Oxford University Press, 2006.

Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory neurons. *Journal of Computational Neuroscience*, 8:183–208, 2000.