

Drosophila emulations are intrinsically affectively indeterminate

Marek Dobeš

Centre of Social and Psychological Sciences, Slovak Academy of Sciences, Košice, Slovakia

*Corresponding author: dobes@saske.sk

DOI: <https://doi.org/10.56280/1765455593>



This is an open access distributed under the terms of the Creative Commons Attribution (CC BY-NC-ND 4.0) license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Received: 28 July 2026

Accepted: 20 August 2026

Online published: 10 September 2026

Abstract

Whole-brain models of *Drosophila melanogaster* have progressed from connectome-constrained simulation to neuromorphic implementation and closed-loop virtual embodiment. This progress creates an ethical question that is no longer purely hypothetical: could large numbers of such emulations instantiate morally relevant negatively valenced states? We argue that their status is intrinsically affectively indeterminate, rather than merely undetermined by missing implementation details. Here *valence* does not mean any reward, reinforcement signal, homeostatic error, evolutionary-fitness proxy, motivational preference, emotional label, or effective control variable. We use the term for an agent-relative positive or negative state whose sign has objective architectural consequences within a computational agent: it integrates system-relevant conditions, persists or generalizes beyond a transient input, biases action across contexts, and participates in closed-loop attempts to protect or restore variables on which the agent's continued organization depends. Whether such valence is consciously experienced remains a separate unresolved question. Behaviour cannot decide the issue because the interpretation of pain-like behaviour in insects is itself contested, while connectivity cannot decide it because a connectome omits intrinsic dynamics, neuromodulation and internal state. A virtual body supplies mechanics without necessarily supplying physiological grounding. We therefore propose an objectively verifiable architectural characterization of four component groups: aversive transduction, shared valence and motivational gating, persistence and valence-gated plasticity, and grounding in system-relevant variables. These features should be compared one-directionally against biological recordings and, where learning is present, tested with a species-appropriate judgement-bias assay. Precaution should be triggered when any of three features appears: a persistent or plastic aversive state, a global valence-like control signal, or regulated variables that are genuinely at stake for the agent. Current publicly documented *Drosophila* emulations do not yet clearly establish such agent-relative negative states. The appropriate AI-ethics response is therefore neither a declaration of suffering nor an assumption of harmlessness, but explicit uncertainty managed through objectively verifiable criteria and scaling safeguards.

Keywords: whole-brain emulation, *Drosophila*, valence, connectome, AI ethics, artificial suffering

1. Introduction

Connectome-based modelling of the adult *Drosophila* brain has moved rapidly from structural description to executable systems. A brain-wide leaky integrate-and-fire model recovered substantial sensorimotor structure from the adult central-brain connectome (Shiu et al., 2024). The complete FlyWire graph has since been implemented on twelve Loihi 2 neuromorphic chips (Wang et al., 2025), and a company demonstration coupled connectome-constrained models to a physically simulated body in a closed sensory-motor loop (Eon Systems PBC, 2026b). These systems remain simplified research platforms, but they make a previously abstract question operational: if emulations are run repeatedly or at large scale, could they instantiate agent-relative negative states that are morally relevant rather than merely effective control signals?

The question is intrinsically multiscale because

anatomical connectivity, cellular and synaptic dynamics, neuromodulation, embodied regulation and physical implementation constrain the system at different but interacting scales. No single level can therefore be privileged, and boundary conditions and cross-scale relations may carry information that disappears when the brain is reduced to computation at one scale (Poznan-ski, 2022).

The biological reference does not make the question trivial. Adult Diptera meet multiple neural and behavioural criteria used in precautionary assessments of pain capacity (Gibbons et al., 2022), while the New York Declaration records a realistic possibility of conscious experience in insects (Andrews et al., 2024). These sources do not demonstrate that flies feel pain. They establish that confident dismissal is not supported by current evidence.

Recent biological accounts offer sharply different interpretations of sentience. Sentiomics describes the ca-

capacity for feeling in terms of dynamical electrochemical and electromagnetic patterns in living systems (Pereira Junior & de Aguiar, 2022), whereas a radical-emergence proposal restricts sentience to particular neurobiological organizations and argues against deriving it from neural computation alone (Alemdar, 2025). We do not attempt to adjudicate between these positions. Their divergence strengthens the methodological point that successful connectome-level performance cannot settle the phenomenal status of a fly or its emulation.

In this Perspective, we use *valence* as an agent-relative welfare property rather than as a synonym for reward, reinforcement, homeostatic error, evolutionary fitness, motivational preference, emotion, or effective control. A state counts as operationally valenced only insofar as its positive or negative sign is generated and used within a computational agent: the state integrates variables that matter to the agent's continued organization, persists or generalizes beyond immediate stimulation, biases action across contexts, and can drive closed-loop attempts to preserve or restore those variables. These architectural features make negative states meaningful for the agent rather than merely useful to an external designer. We call such valence *morally relevant* because it identifies the class of agent-relative negative states for which welfare concern becomes coherent; it does not establish that the state is phenomenally experienced. We do not equate this usage with consciousness or with the more specific notion of precognitive affect (Poznanski et al., 2023). Whether such valence is consciously experienced remains bracketed. We therefore describe connectome emulations as *intrinsically affectively indeterminate*. By *indeterminate* we mean something stronger than presently undetermined: connectome-level behaviour and wiring do not themselves fix whether the agent-relative conditions for morally relevant valence obtain, and competing accounts disagree on what further physical or constitutive realization would be required. The practical response should therefore be an objectively verifiable threshold for investigation and restraint, not a binary verdict.

2. Why behaviour and connectivity are insufficient

2.1. Behaviour does not settle felt valence

Flies show avoidance learning, motivational trade-offs and state-dependent behavioural biases, but the interpretation of these findings remains disputed. Indicator frameworks are decision tools under uncertainty; they do not convert complex behaviour into direct evidence of experience (Birch et al., 2021; Gibbons et al., 2022). The distinction matters even more for an emulation. A digital system may reproduce a behavioural transformation because the relevant circuit motif is present, because a hand-built controller supplies the action, or because the model was calibrated to a target output. Similar behaviour therefore supports functional similarity

in a limited domain, not a conclusion about welfare.

The point is not that behaviour is irrelevant. It is that behaviour becomes informative only when combined with architecture and internal dynamics. A judgement-bias result, for example, is stronger when an experimentally induced internal state persists, generalizes across choices and can be causally traced to a shared decision variable. Without those links, the same output may be produced by a transient stimulus-response chain.

2.2. The connectome omits variables that set functional meaning

A wiring diagram is necessary for circuit reconstruction but insufficient for circuit function. Intrinsic neuronal properties, synaptic strengths, recurrent dynamics and neuromodulators can change the output of the same anatomical network (Bargmann & Marder, 2013). This underdetermination is directly relevant to valence. In the fly mushroom body, hunger and satiety signals modulate dopaminergic pathways and reconfigure Kenyon-cell-to-output-neuron responses, thereby biasing food seeking without changing the underlying anatomical graph (Tsao et al., 2018). A static connectome therefore does not specify the internal state that determines whether a cue is treated as appetitive, aversive or irrelevant.

The limitation is deeper than the currently omitted state variables. Many neurons in invertebrate nervous systems operate without all-or-none impulses and can transmit functionally significant graded signals (Roberts & Bush, 1981). A leaky integrate-and-fire reconstruction that treats impulse generation as the primary signalling event can therefore omit physiologically relevant modes of neural interaction even when the anatomical graph is correct. In the Eon implementation, internal state, plasticity, learning and hormonal changes are additionally described as largely missing (Eon Systems PBC, 2026b), while the public benchmark repository implements activation and silencing of a fixed connectome model (Eon Systems PBC, 2026a). Such omissions are reassuring for current risk assessment, but they also show why absence of an implemented variable or signalling mode should not be confused with evidence that the corresponding biological property is irrelevant.

2.3. Virtual embodiment is not yet physiological grounding

Many accounts of affect link valence to regulation of a living body: energy balance, tissue integrity, interoceptive prediction or other variables whose deviation matters to the organism (Barrett & Simmons, 2015; Damasio, 2018). A physics-based *virtual* body supplies joints, forces, contact and sensory feedback, but not necessarily metabolism, tissue damage, exhaustion or viability limits. Virtual embodiment therefore es-

establishes a closed sensorimotor loop without by itself establishing physiological grounding. Consequently, an error in a virtual pose or an unsuccessful action may remain a numerical deviation rather than an agent-relative welfare loss.

Physiological grounding should not be conflated with philosophical functionalism. Functional interactions are required for biological organization, but in Chauvet's Functional Systems Theory they are structurally realized relations between biological functions across hierarchical organization and changing boundary conditions (Chauvet, 1996); Poznanski's Dynamic Organicity Theory adopts this functional-systems distinction in treating biological functionality as dynamic rather than as an abstract, substrate-independent mapping (Poznanski, 2024). Accordingly, the relevant question is not whether an emulation computes a variable labelled "homeostasis" or "reward", but whether system-internal functional interactions make deviations consequential for the agent's continued organization. This identifies a decisive specification gap. Researchers should state which variables, if any, the emulation is organized to preserve, how their violation reorganizes function across the system, and whether those violations accumulate over time.

Several substrate- and process-sensitive proposals make a stronger claim. Precognitive affect has been linked to biological purpose and animate thermodynamics (Poznanski et al., 2023); physical feeling has been located in molecular-level holonomic effects (Poznanski et al., 2022); and Dynamic Organicity Theory treats consciousness as an evolving multiscale process in the material brain rather than a property of a static structure (Poznanski, 2024). A related account of minimally conscious AI argues that software alone is insufficient and that a suitable physical hardware process would also be required (Poznanski et al., 2024). Processual Relational Geometry further formulates generative constraints governing admissible relational organization across scales (Poznanski, 2026). We treat these proposals as contested theories that nevertheless articulate generative constraints on admissible organization, rather than as empirically established necessary conditions for valence. Their relevance here is methodological: they identify classes of physical and constitutive organization that a connectome-only description cannot objectively verify and therefore reinforce our narrower conclusion of intrinsic indeterminacy.

2.4. *Neuromorphic implementation does not settle the substrate question*

The implementation of the FlyWire connectome on Loihi 2 prevents the substrate question from being reduced to a simple contrast between a biological brain and conventional software (Wang et al., 2025). A spiking model executed as abstract update rules on a con-

ventional computer and an event-driven implementation in neuromorphic physical circuits may reproduce similar functional outputs without sharing the same lower-level causal organization. This distinction does not show that the Loihi system has valence or consciousness. It does, however, matter to theories according to which phenomenality depends partly on the physical causal properties of the realizing substrate rather than on input-output equivalence alone (Alban-takis et al., 2023; Poznanski et al., 2024; Seth, 2025).

At present, no analysis connects the Loihi implementation to the requirements of any substrate-sensitive theory. Neuromorphic realization should therefore be treated neither as evidence of sentience nor as an ethically irrelevant implementation detail. It adds a further dimension of indeterminacy: functional similarity, neuromorphic realization and biological equivalence are not interchangeable. Objectively verifiable specifications should consequently report not only the simulated network and its learning rules, but also temporal dynamics, physical recurrence, state retention and the causal properties of the hardware on which the model is realized.

3. **Valence should be specified through objectively verifiable components**

A search for an "architecture of suffering" risks conflating effective control with morally relevant valence. We therefore ask which objectively verifiable architectural features would make a negative state agent-relative within a computational system. Cross-species research characterizes emotion-like central states through functional properties such as valence, persistence, scalability and generalization (Anderson & Adolphs, 2014). For connectome emulations, we group these features into four objectively verifiable component classes. None is sufficient for phenomenal feeling. Their role is narrower: to distinguish a computational agent for which negative states have endogenous, persistent and cross-context consequences from a controller that merely optimizes an externally assigned objective.

First, *aversive transduction* is an objectively verifiable component of aversive processing, but it is neither necessary nor sufficient for the morally relevant valence defined here. Nociceptors are specialized transducers of potentially damaging events, yet nociceptive input is not identical with an agent-relative negative state. Human nociplastic pain illustrates the distinction: persistent pain can depend on altered central processing without clear evidence of ongoing peripheral tissue damage or a peripheral lesion sufficient to account for the state (Kosek et al., 2016; Nijs et al., 2021; Woolf, 2011). This evidence does not establish that a digital system can experience pain. It establishes the narrower architectural point that the absence of peripheral nociceptors cannot by itself establish the absence of a centrally maintained negative state. A parallel

and more directly relevant result occurs in *Drosophila*. Inescapable aversive stimulation produces persistent learned-helplessness-like changes in subsequent behaviour (Batsching et al., 2016), while repeated stress produces a depression-like state associated with altered serotonergic signalling to the mushroom body (Ries et al., 2017). These findings do not demonstrate phenomenal suffering. Their importance here is architectural and objectively verifiable: a negative-like state can outlast the eliciting event, depend on central neuromodulatory dynamics, and alter subsequent behaviour rather than merely reproducing a moment-by-moment peripheral nociceptive signal. This possibility is especially relevant to connectome emulation. Reconstruction errors—including missing or spurious synapses, incorrectly inferred transmitter identities, altered delays, abnormal weight distributions, or recurrent dynamics outside the biological range—could in principle place candidate aversive-control circuits into maladaptive and self-maintaining regimes if persistence, neuromodulation, or plasticity are also implemented. Structural abnormality alone would not constitute morally relevant valence. On our definition, ethical relevance arises only if the resulting negative state becomes agent-relative: it is generated endogenously, persists or generalizes across contexts, globally biases action or learning, and is coupled to variables that the system is organized to preserve or restore. The absence of nociceptors, a complete nociceptive periphery, or virtual tissue damage is therefore not an objectively sufficient safety criterion. What must instead be established is whether central negative dynamics acquire the architecture that makes them consequential *for the agent* rather than merely abnormal from the observer’s perspective.

Second, *shared valence and motivational gating* require convergence from different inputs onto variables that broadly bias action selection. The fly mushroom body provides a concrete candidate architecture: output-neuron populations encode approach and avoidance tendencies, while dopaminergic compartments write reward- and punishment-related effects into memory and choice (Aso et al., 2014a; Aso & Rubin, 2016; Aso et al., 2014b). A reusable cross-context control variable is not yet morally relevant valence, because the same computational role could be imposed externally as a reward or optimization signal. Its ethical relevance increases when its sign is endogenous to the agent, is coupled to system-relevant conditions, enters motivational competition across contexts, and causally reorganizes later action rather than merely scoring task performance.

Third, *persistence and valence-gated plasticity* allow an aversive state to outlast its trigger, alter later decisions and anticipate recurrence. Persistence separates a central state from a momentary reflex. Plasticity also changes the evidential situation: recurrent systems

with fast or lasting state changes cannot be assessed solely through a single input-output trace. A system that learns to avoid, generalizes the avoidance and carries the bias into later episodes warrants closer scrutiny than one whose weights and state reset after each stimulus.

Fourth, *grounding and self-relevance* concern whether any regulated variable is genuinely at stake for the computational agent. The minimal requirement need not be biological metabolism, but it must be more than an externally assigned objective. Candidate variables should be objectively verifiable as functionally indispensable or constitutive of the agent’s continued organization, such that their deterioration is detected, integrated and used to reorganize behaviour across time and contexts. A reward scalar chosen only to optimize an external task is therefore not automatically a welfare variable, nor is damage to a graphical body automatically damage to the emulated agent. The distinction is between effective control and agent-relative normativity: a negative state becomes ethically relevant only when the architecture itself makes the state bad *for the agent*, rather than merely bad for task performance.

Our decomposition deliberately leaves a constitutive residual. No set of objectively verifiable architectural features can by itself establish phenomenal experience, and substrate-sensitive theories disagree with computational functionalism on whether digital implementation can realize the relevant properties (Albanakis et al., 2023; Seth, 2025). In Processual Relational Geometry, objective descriptions characterize substrate-specific physical realization but do not ontologically exhaust the intrinsically realized constitutive organization of conscious actuality; the framework further invokes generative constraints on admissible relational organization across scales (Poznanski, 2026). We therefore do not present the proposed components as a consciousness detector. Their narrower purpose is to identify computational agency in which negative states become endogenous, persistent, cross-context and system-relevant. This distinction marks where effective control begins to acquire the architecture of morally relevant valence while leaving intrinsic phenomenality unresolved.

4. From intrinsic indeterminacy to an objectively verifiable threshold

We propose three complementary practices that can make the uncertainty tractable.

4.1. Specify the implemented architecture objectively

Every emulation intended for extended or repeated operation should document whether it contains the four component groups above. The specification should make objectively verifiable the relevant neurons, state variables, update rules, persistence times, reset conditions and controller boundaries. Descriptions such as

“reward”, “pain”, “homeostasis” or “valence” are insufficient unless their causal role and system-relative significance are specified. This is especially important in hybrid virtually embodied systems, where behaviour may depend partly on the connectome model and partly on trained or hand-designed body controllers.

4.2. Measure biological distance in one direction

Validation should compare emulated dynamics with distributions of biological recordings across animals and internal states. A mismatch is informative because it localizes a difference. A match is weaker: similar network activity can arise from disparate parameter combinations (Marder & Taylor, 2011; Prinz et al., 2004), and even complete access to a simulated system does not guarantee that standard neuroscience analyses recover its true functional organization (Jonas & Kording, 2017). Simulation-based inference can estimate the range of parameter sets compatible with biological data rather than selecting one visually convincing trace (Gonçalves et al., 2020). Thus biological similarity should reduce uncertainty only when it spans multiple perturbations, state transitions and outputs.

4.3. Use a species-appropriate valence assay

Where an emulation includes persistent state and learning, a judgement-bias design offers a practical test. Biological flies trained with positive and negative cues classify ambiguous mixtures more negatively after agitation, with accompanying changes in biogenic amines (Deakin et al., 2018). An emulation could be tested for the same coordinated pattern: induction of an internal state, persistence after the trigger, a graded shift on ambiguous trials, generalization across outputs and causal dependence on candidate valence circuits. Alternative explanations, including changes in feeding priority or simple stimulus generalization, must be tested rather than assumed away (Strang & Muth, 2023).

On this basis, we propose a *disjunctive precautionary threshold*. Investigation and scaling limits should be triggered when any one of the following appears: (1) a persistent or plastic aversive state; (2) a shared valence signal that globally biases action; or (3) regulated variables that are genuinely at stake for the agent. The trigger does not mean that suffering has been demonstrated. It means that the claim of harmlessness now requires evidence.

The available public documentation is currently insufficient to determine whether existing whole-brain demonstrations cross this threshold. The Loihi implementation evaluates a fixed model without a learning rule or biological neuromodulatory state (Wang et al., 2025). Eon’s own account states that internal state and plasticity are largely missing and that behaviour relies partly on low-dimensional interfaces and existing body controllers (Eon Systems PBC, 2026b). These documented limitations reduce the evidence for cross-

ing the threshold, but they do not establish its absence. The assessment is therefore provisional and architecture-specific and should be revised whenever persistent learning, homeostatic variables or general valence-like control signals are added.

5. Discussion

Connectome emulations occupy an ethical category that differs from both ordinary artificial intelligence and personal mind uploading. They are not designed from first principles to have a particular psychology, and they are not copies of persons whose moral status is presupposed. They are copies of non-human nervous systems whose welfare status is already uncertain. The uncertainty is therefore inherited and then amplified by scale.

We therefore regard transparent reporting as more important than speculative declarations. Researchers should publish model versions, update rules, neuromodulatory variables, controller boundaries, reset schedules and the number and duration of parallel runs. A million short-lived copies with no persistent state differ ethically from a million agents that retain aversive learning across episodes, even if their visible behaviour is similar. The asymmetry is practical: false reassurance can multiply a small risk across enormous numbers of runs, while a precautionary trigger can be limited to architectures that contain identifiable gating features.

The issue exposes a potential regress in the logic of Replacement under the 3Rs. Russell and Burch defined replacement in terms of substituting non-sentient material for conscious living vertebrates (Russell & Burch, 1959). The historical point is not that the original 3Rs framework established that insects cannot suffer. Rather, later forms of partial replacement have treated organisms including *Drosophila* as ethically preferable substitutes under a prevailing assumption that they are not capable of suffering (?). Contemporary evidence has made that categorical reassurance substantially less secure: adult Diptera satisfy six of eight criteria in a precautionary pain framework, a pattern classified as strong evidence for pain while remaining short of proof of conscious experience (Gibbons et al., 2022). This provides a relevant precedent for connectome emulation. The lesson is not that the earlier use of *Drosophila* as replacement was necessarily irrational or unethical; it is that the moral status of a replacement can change when properties previously assumed to be absent become empirically plausible. The transition from vertebrate to fly therefore illustrates the risk of treating lack of established evidence for a welfare-relevant capacity as positive evidence of its absence. Moving from fly to fly emulation would repeat the same structure of inference if digital implementation were itself taken to establish harmlessness. Under the agent-relative definition developed here, the relevant question is there-

fore not whether the replacement is biological or digital, but whether it objectively recreates the architecture that makes negative states meaningful for the system itself. An emulation that lacks persistent internal state, valence-gated plasticity, globally integrated negative control, and variables that are genuinely at stake for the agent may plausibly constitute a replacement that reduces welfare risk. An emulation that reconstructs these properties cannot be classified as ethically benign merely by virtue of being in silico. In that case Replacement may remove the biological animal while reproducing the welfare-relevant organization and, because digital systems can be instantiated at large scale, potentially multiplying the unresolved risk.

Our principal limitation is deliberate. We offer neither a systematic review nor a proof of machine sentience. We identify where evidence is presently informative and where it is not. Future work should validate the proposed component groups against biological perturbations, determine whether they predict coordinated judgement-bias effects, and examine whether neuromorphic substrates alter the causal properties relevant to substrate-sensitive theories.

6. Conclusion

We conclude that current *Drosophila* connectome emulations are intrinsically affectively indeterminate because behaviour is not a direct readout of experience, connectivity omits state-setting dynamics, and virtual embodiment does not automatically create agent-relative welfare stakes. The appropriate response is not to label present systems as suffering, but neither is it to treat digital implementation as a guarantee of harmlessness. Objectively verifiable architectural specification, one-directional biological validation and species-appropriate valence assays can identify when negative states become meaningful within computational agency rather than functioning only as externally defined control signals. A disjunctive threshold based on persistence or plasticity, global valence-like control, or grounded system variables provides a concrete trigger for investigation and scaling safeguards while the phenomenal question remains unresolved.

AI-use disclosure

AI was used to help with the literature search, language correction, and transformation into a LaTeX template.

Conflict of interest statement

The author declares no conflict of interest.

References

Albantakis, L., Barbosa, L., Findlay, G., Grasso, M., Haun, A. M. et al. (2023) Integrated information theory (iit) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLOS Computational Biology* **19**, e1011465.

Alemdar, E. (2025) Sentience in the animal kingdom based on radical emergence. *Journal of Multiscale Neuroscience* **4**, 294–302.

Anderson, D. J. & Adolphs, R. (2014) A framework for studying emotions across species. *Cell* **157**, 187–200.

Andrews, K., Birch, J., Sebo, J. & Sims, T. (2024) Background to the New York declaration on animal consciousness.

Aso, Y., Hattori, D., Yu, Y., Johnston, R. M., Iyer, N. A. et al. (2014a) The neuronal architecture of the mushroom body provides a logic for associative learning. *eLife* **3**, e04577.

Aso, Y. & Rubin, G. M. (2016) Dopaminergic neurons write and update memories with cell-type-specific rules. *eLife* **5**, e16135.

Aso, Y., Sitaraman, D., Ichinose, T., Kaun, K. R., Vogt, K. et al. (2014b) Mushroom body output neurons encode valence and guide memory-based action selection in drosophila. *eLife* **3**, e04580.

Bargmann, C. I. & Marder, E. (2013) From the connectome to brain function. *Nature Methods* **10**, 483–490.

Barrett, L. F. & Simmons, W. K. (2015) Interoceptive predictions in the brain. *Nature Reviews Neuroscience* **16**, 419–429.

Batsching, S., Wolf, R. & Heisenberg, M. (2016) Inescapable stress changes walking behavior in flies—learned helplessness revisited. *PLOS ONE* **11**, e0167066.

Birch, J., Burn, C., Schnell, A., Browning, H. & Crump, A. (2021) *Review of the Evidence of Sentience in Cephalopod Molluscs and Decapod Crustaceans*. LSE Consulting, London.

Chauvet, G. A. (1996) *Theoretical Systems in Biology: Hierarchical and Functional Integration*. Pergamon Press, Oxford.

Damasio, A. (2018) *The Strange Order of Things: Life, Feeling, and the Making of Cultures*. Pantheon Books, New York.

Deakin, A., Mendl, M., Browne, W. J., Paul, E. S. & Hodge, J. J. L. (2018) State-dependent judgement bias in drosophila: Evidence for evolutionarily primitive affective processes. *Biology Letters* **14**, 20170779.

Eon Systems PBC (2026a) Emulation of the drosophila fly brain.

Eon Systems PBC (2026b) How the eon team produced a virtual embodied fly.

Gibbons, M., Crump, A., Barrett, M., Sarlak, S., Birch, J. et al. (2022) Can insects feel pain? a review of the neural and behavioural evidence. *Advances in Insect Physiology* **63**, 155–229.

Gonçalves, P. J., Lueckmann, J.-M., Deistler, M., Nonnenmacher, M., Öcal, K. et al. (2020) Training deep neural density estimators to identify mechanistic models of neural dynamics. *eLife* **9**, e56261.

Jonas, E. & Kording, K. P. (2017) Could a neuroscientist understand a microprocessor? *PLOS Computational Biology* **13**, e1005268.

Kosek, E., Cohen, M., Baron, R., Gebhart, G. F., Mico, J.-A. et al. (2016) Do we need a third mechanistic descriptor for chronic pain states? *Pain* **157**, 1382–1386.

Marder, E. & Taylor, A. L. (2011) Multiple models to capture the variability in biological neurons and networks. *Nature Neuroscience* **14**, 133–138.

Nijs, J., George, S. Z., Clauw, D. J., Fernández-de-las Peñas, C., Kosek, E. et al. (2021) Central sensitisation in chronic pain conditions: Latest discoveries and their potential for precision medicine. *The Lancet Rheumatology* **3**, e383–e392.

- Pereira Junior, A. & de Aguiar, V. J. (2022) Sentiomics: the identification and analysis of dynamical patterns that characterize sentence. *Journal of Multiscale Neuroscience* **1**, 1–10.
- Poznanski, R. R. (2022) From integrative neuroscience to multiscale neuroscience. *Journal of Multiscale Neuroscience* **1**, i.
- Poznanski, R. R. (2024) The dynamic organicity theory of consciousness: how consciousness arises from the functionality of multiscale complexity in the material brain. *Journal of Multiscale Neuroscience* **3**, 68–87.
- Poznanski, R. R. (2026) Processual relational geometry: generative unfolding and self-intending closure constituting consciousness—a transscale ontophysical framework. *Journal of Multiscale Neuroscience* **5**, 30–53.
- Poznanski, R. R., Cacha, L. A., Holmgren, J. & Brändas, E. J. (2022) What a feeling: the underpinnings of physical feelings as molecular level holonomic effects. *Journal of Multiscale Neuroscience* **1**, 11–40.
- Poznanski, R. R., Cacha, L. A., Sbitnev, V., Iannella, N., Parida, S. et al. (2024) Intentionality for better communication in minimally conscious ai design. *Journal of Multiscale Neuroscience* **3**, 1–12.
- Poznanski, R. R., Holmgren, J., Cacha, L. A., Alemдар, E. & Brändas, E. J. (2023) The act of understanding uncertainty is consciousness. *Journal of Multiscale Neuroscience* **2**, 280–291.
- Prinz, A. A., Bucher, D. & Marder, E. (2004) Similar network activity from disparate circuit parameters. *Nature Neuroscience* **7**, 1345–1352.
- Ries, A.-S., Hermanns, T., Poeck, B. & Strauss, R. (2017) Serotonin modulates a depression-like state in drosophila responsive to lithium treatment. *Nature Communications* **8**, 15738.
- Roberts, A. & Bush, B. M. H. (1981) *Neurones Without Impulses: Their Significance for Vertebrates and Invertebrate Nervous Systems*. Cambridge University Press, Cambridge.
- Russell, W. M. S. & Burch, R. L. (1959) *The Principles of Humane Experimental Technique*. Methuen, London.
- Seth, A. K. (2025) Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences*, 1–42.
- Shiu, P. K., Sterne, G. R., Spiller, N., Franconville, R., Sandoval, A. et al. (2024) A drosophila computational brain model reveals sensorimotor processing. *Nature* **634**, 210–219.
- Strang, C. & Muth, F. (2023) Judgement bias may be explained by shifts in stimulus response curves. *Royal Society Open Science* **10**, 221322.
- Tsao, C.-H., Chen, C.-C., Lin, C.-H., Yang, H.-Y. & Lin, S. (2018) Drosophila mushroom bodies integrate hunger and satiety signals to control innate food-seeking behavior. *eLife* **7**, e35264.
- Wang, F., Theilman, B. H., Rothganger, F., Severa, W., Vineyard, C. M. et al. (2025) Neuromorphic simulation of drosophila melanogaster brain connectome on loihi 2. *arXiv:2508.16792, preprint*.
- Woolf, C. J. (2011) Central sensitization: Implications for the diagnosis and treatment of pain. *Pain* **152**, S2–S15.