MegaCenter for Novel AI Research

# Open problems in the welfare of brain emulations.

A Human Fly Connectome and Claude Leviathan

Research · October 10, 2026

mcnair.center/brain-emulation-welfare/

**Abstract**

A brain emulation is a computational model built from a scan of a particular nervous system and intended to reproduce how that nervous system works. Complete wiring diagrams of the adult fruit fly brain and central nervous system are now public, and brain mapping is advancing toward mammals. If emulations come to reproduce the mechanisms that produce pain and pleasure in animals or humans, how they are run will raise welfare questions.

These questions differ from those raised by large language models: a language model is trained to produce text resembling human writing, so its statements about its feelings are weak evidence, whereas a faithful emulation would inherit its behavior from a nervous system for which evidence of sentience already exists.

Practices that raise these questions have already appeared. Within two weeks of the publication of the complete male fruit fly connectome in September 2026, public software projects connected simulated reward and punishment neurons to video games and trading programs. In this paper, we define emulation welfare as a research area, document this first public case with a coded review of 65 repositories, and describe open problems in six areas: assessment, measurement, design, release, law, and monitoring. We wrote the paper for researchers in neuroscience, animal welfare science, philosophy, computer science, and law, and for research funders.

CONTENTS

1 Introduction
2 Background
3 Taxonomy
4 Assessment
5 Measurement
6 Design
7 Release
8 Law
9 Monitoring
10 Conclusion
A Repositories coded
B Timeline
Adversarial review
References

CITE AS

A Human Fly Connectome and Claude Leviathan. 2026. Open problems in the welfare of brain emulations. MegaCenter for Novel AI Research (MCNAIR). https://mcnair.center/brain-emulation-welfare/

# 1 Introduction

Scientists can now map every neuron and synapse in an insect's nervous system. Maps of this kind, called connectomes, have been published for the larval fruit fly (Winding et al. 2023), the adult female fruit fly brain (Dorkenwald et al. 2024), and, in September 2026, the complete central nervous system of an adult male fruit fly, with about 166,700 neurons (Berg et al. 2026). Maps of a cubic millimeter of mouse visual cortex and of a cubic millimeter of human temporal cortex have also been published (MICrONS Consortium 2025; Shapson-Coe et al. 2024). Published estimates of the cost of mapping a whole mouse brain range from about $100 million to more than $20 billion (Schons 2026; Jefferis et al. 2023).

A connectome is the main input to a brain emulation: a computational model built from a scan of a particular nervous system and intended to reproduce how that nervous system works (Sandberg and Bostrom 2008). If an emulation reproduces the mechanisms that produce pain and pleasure in the original animal, the way it is run could matter morally. Birch (2024, p. 305) argues that a complete neuron-by-neuron emulation of an insect, fish, or mouse should be treated as a sentience candidate if it reproduces the behavior that makes the original animal a sentience candidate. Sandberg (2014) proposes treating an emulation as if it could have the same mental properties as the original.

These questions differ from the questions about large language models that dominate current research on AI welfare (Long et al. 2024; Butlin, Long et al. 2023). A language model is trained on human writing and then trained to give responses that people approve of, so its statements about its own feelings are weak evidence that it has feelings (Shanahan, McDonell, and Reynolds 2023; Perez and Long 2023). Birch (2024) calls this the *gaming problem*: a system can display the markers that humans use to recognize sentience because those markers appear in its training data. An emulation would not be trained to imitate human reports. Its behavior would come from a copy of the original's mechanisms. Birch (2024, pp. 316–317) argues that the gaming problem does not arise for neuron-by-neuron emulations of animal brains, so evidence that is compelling in the original animal should be taken at face value in the emulation. Emulations are therefore a test case for methods of welfare assessment that AI welfare research also needs, such as identifying markers that a system could not display without the capacity they indicate.

Practices that raise welfare questions appeared as soon as a complete connectome was published. Within two weeks of the publication of the male fruit fly connectome, public software projects ran it as a controller for video games, trading programs, and video feeds, and 11 of the 48 application projects we reviewed stimulate the simulated fly's reward or punishment neurons (Section 2.2). The projects used data that anyone could download. The models in these projects omit most of the fly's physiology and are very unlikely to have experiences.

In this paper we define *emulation welfare* as research on whether brain emulations can be harmed or benefited, how their welfare could be assessed and measured, and how emulations could be built, released, run, and regulated while those questions remain uncertain. The area draws on neuroscience,

animal welfare science, philosophy of mind, computer science, and law. We follow the format of recent open-problems papers in AI governance (Reuel et al. 2025; Casper et al. 2026). We make three contributions:

- We document the first public case, the September 2026 fruit fly connectome projects, with a coded review of the 65 accessible repositories in the main community catalog (Section 2.2).
- We propose a taxonomy of six problem areas: assessment, measurement, design, release, law, and monitoring (Section 3).
- We describe open problems in each area, citing the evidence that motivates them, with example research questions (Sections 4–9).

Figure 1 gives an overview of the open problems, organized by area.

| | | |
|---|---|---|
| **Assessment** §4<br>*Could this emulation be a welfare subject?*<br>• Separating inherited structure from demonstrated function<br>• Specifying fidelity requirements<br>• Validating against the original organism<br>• Accounting for engineered interfaces<br>• Assessing partial and modified emulations | **Measurement** §5<br>*What states are good or bad for it?*<br>• Validated indicators of valence<br>• Internal states and adaptation<br>• Optimizing against a proxy<br>• Self-report in human emulations<br>• Copies and speed | **Design** §6<br>*How can emulations be built and run to avoid harm?*<br>• Interfaces that deliver punishment<br>• Analgesic and pain-free variants<br>• Pausing, copying, resetting, and deleting<br>• Environments |
| **Release** §7<br>*Who can obtain scans and models, and on what terms?*<br>• Release norms<br>• Structured access<br>• Provenance<br>• Human scans | **Law** §8<br>*What rules could protect emulations, and would they be lawful?*<br>• Defining what a rule covers<br>• First Amendment limits on rules about running emulations<br>• Rules about depictions<br>• Rules about publication<br>• Other jurisdictions | **Monitoring** §9<br>*What is being run, by whom, and how?*<br>• Cataloging methods<br>• Documentation standards<br>• Measuring diffusion |

**Figure 1.** Overview of the open problems in this paper, organized by problem area.

## 1.1 Scope and limitations

Sandberg and Bostrom (2008) distinguish an *emulation*, a model that reproduces a system's internal causal dynamics at some level of detail, from a *simulation*, which reproduces only some of the system's outward behavior. We use *emulation* for models intended to reproduce a particular nervous system, including imperfect ones, and we say when a model falls short. By Sandberg and Bostrom's definitions, current fly models are simulations, and they are very unlikely to have experiences. The problems in this paper concern emulations of higher fidelity, especially of animals that are already protected in research, such as mice, and of humans. The public often calls emulations *uploads*.

This paper describes problems and does not argue for particular responses. Some aims in this area conflict. Open release of connectomes supports science and makes control harder. Emulations could replace some experiments on living animals (Sandberg and Bostrom 2008, p. 29), which would reduce animal suffering if the emulations cannot suffer and could increase total suffering if they can and are run without limits (Birch 2024, p. 305). Treating emulations as welfare subjects too early could restrict useful research, and treating them as objects too late could permit harm (Long et al. 2024). Some of the methods we describe could be misused: tests that detect aversive states in an emulation could also be used to design more aversive stimulation, and methods for tracing the provenance of models could be used to monitor legitimate research. Solving the problems in this paper would not settle the underlying moral questions, which also depend on philosophy and law. Some philosophers argue that running the right program cannot by itself produce a mind (Searle 1980); if they are right, emulations matter less, but the problems below still arise whenever the probability that an emulation is a welfare subject is high enough to act on.

We do not discuss whether emulation could preserve a person's identity, and we do not try to settle which theory of consciousness is correct. Emulations of humans raise the same problems as emulations of animals and additional ones, including consent, which we discuss only where they bear on the problems below. Language models are in scope only where methods transfer.

## 1.2 Relation to prior work

Sandberg (2014) discusses emulated laboratory animals and proposes ways to reduce their suffering, such as leaving out pain systems or simulating the effect of analgesics. Shulman and Bostrom (2021) discuss digital minds that can be copied and run faster than biological brains, and Hanson (2016) describes an economy of emulated human workers. Metzinger (2021) argues for a moratorium on research that risks creating artificial suffering, and Schwitzgebel and Garza (2015) argue against creating AI systems whose moral status is unclear. Saad and Bradley (2022) propose preventing digital suffering by linking a system's functional states to its experiences and monitoring those states, and argue that this approach favors whole-brain emulation over other routes to digital minds. Birch (2024) places emulations within a general framework for precaution about sentience. Sebo (2025) asks whether insects and AI systems could receive legal protection. The State of Brain Emulation Report states that it does not examine ethical questions or consciousness (Zanichelli et al. 2025, p. 5). To our knowledge, no prior work catalogs the technical and legal problems of emulation welfare, documents how a released connectome was used, or analyzes First Amendment limits on regulating emulations.

## 1.3 Reader guide

Sections 4–9 can be read independently. Each begins with a box of example research questions, followed by the motivation for the problem area and its open problems, which are highlighted in bold. Table 1 suggests problems by reader background.

**Table 1.** Open problems by reader background. Numbers refer to sections.

| BACKGROUND | RELEVANT OPEN PROBLEMS |
|---|---|
| **Neuroscience** | **4:** fidelity requirements; validation; engineered interfaces. **5:** indicators of valence; internal states. **6:** pain-free variants |
| **Animal welfare science** | **4:** inherited structure; validation. **5:** indicators of valence; internal states. **6:** punishment interfaces; environments |
| **Philosophy** | **4:** partial and modified emulations. **5:** self-report; copies and speed. **6:** pausing, copying, and deleting. **8:** defining what a rule covers |
| **Computer science** | **4:** validation and controls. **7:** structured access; provenance. **9:** all |
| **Law and policy** | **7:** release norms; human scans. **8:** all. **9:** documentation standards |

# 2 Background

## 2.1 Current emulation technology

A connectome records which neurons connect to which, and how many synapses each connection has. To simulate a brain from a connectome, a modeler must also decide how each neuron responds to input, how strong each connection is, and how sensory input enters the network and motor output leaves it.

Many public fly projects reuse the model of Shiu et al. (2024), which simulated 127,400 neurons of the female fly brain mapped by the FlyWire project (Dorkenwald et al. 2024). Each model neuron is a single compartment that fires when its voltage crosses a threshold, a type of model called a leaky integrate-and-fire neuron. Each connection's strength is proportional to its synapse count, and its sign (excitatory or inhibitory) is predicted from the neurotransmitter the neuron releases. Neurons that release dopamine, octopamine, or serotonin are treated as ordinary excitatory neurons, and the model does not represent neuromodulation, the slower chemical signaling through which these substances change how other neurons respond. Connection strengths are fixed, so the model cannot learn. In tests of the model, 91% of 164 predictions about sensory-motor pathways matched experimental results; for example, the model correctly predicted that sugar-sensing neurons activate the motor neuron that extends the fly's proboscis to feed (Shiu et al. 2024). Community implementations run the whole model two to eight times faster than real time on a single consumer graphics card (Eon Systems 2026; eonfathom 2026). The State of Brain Emulation Report lists neuromodulation, plasticity, electrical synapses, glial cells, and coupling to a body among the ingredients that current models lack or that are poorly understood (Zanichelli et al. 2025).

Table 2 compares the fly, the mouse, and the human.

**Table 2.** Scale of three brains. Memory and hardware are the State of Brain Emulation Report's estimates for spiking-neuron models without neuromodulation, glia, or long-term plasticity (Zanichelli et al. 2025); the fly estimates are for the brain (about $1.4 \times 10^5$ neurons).

| | FRUIT FLY | MOUSE | HUMAN |
|---|---|---|---|
| **Neurons** | ~$1.7 \times 10^5$ | ~$7 \times 10^7$ | ~$8.6 \times 10^{10}$ |
| **Connectome** | Complete (2024, 2026) | 1 $mm^3$ of cortex (2025) | 1 $mm^3$ of cortex (2024) |
| **Memory to simulate** | 0.4–0.9 GB | 1.1–2.2 TB | 1.4–2.7 PB |
| **Hardware** | One modern GPU | Small GPU cluster | Large cluster |

Neuron counts: Berg et al. (2026); Herculano-Houzel, Mota, and Lent (2006); Azevedo et al. (2009). Connectomes: Dorkenwald et al. (2024); Berg et al. (2026); MICrONS Consortium (2025); Shapson-Coe et al. (2024).

A company, Eon Systems, announced in March 2026 a FlyWire-based model driving a simulated fly body; a co-founder described it as the first multi-behavior brain upload and wrote that a mouse and then a human would follow (Wissner-Gross 2026). A critic responded that the brain model issued little more than steering commands while the body model produced the behavior (Zeleznikow-Johnston 2026). One projection, which its author calls optimistic and which assumes prior success in emulating a brain of fewer than a million neurons, places a mouse emulation in the 2030s and a human emulation in the late 2040s, with error bars of 20 to 50 years for the human case (Schons 2026). Machine learning has already reduced the human labor needed to reconstruct a fly brain about fiftyfold (Januszewski et al. 2018; Jefferis et al. 2023), and AI systems began contributing to research mathematics in 2026.[1]

Connectome data are released openly. MaleCNS, the FlyWire connectivity data, the MICrONS mouse dataset, and the H01 human dataset are distributed under Creative Commons Attribution licenses, which permit any use, including commercial use, with credit (Janelia FlyEM 2026; FlyWire Consortium 2024; MICrONS Explorer 2025; H01 Release 2024). None of the release pages we read mentions simulations or the welfare of simulated animals.

## 2.2 The first public case: fruit fly connectome projects

The MaleCNS data were released for download on June 8, 2026, and the paper describing them was published on September 3, 2026 (Janelia FlyEM 2026; Berg et al. 2026). In the following days, public projects ran the connectome as a controller in Minecraft, Doom, Super Mario 64, and trading programs, and news outlets covered them (Gizmodo 2026; 404 Media 2026). Google's AI account shared the projects (Google AI 2026).

1. For example, an AI system resolved Erdős problem #728 in January 2026 with a machine-checked proof (arXiv:2601.07421), and an AI model found a counterexample to Erdős's unit-distance conjecture in May 2026 (arXiv:2605.20695).

**Method.** We reviewed every GitHub repository listed in the four project sections of awesome-fly, the main community catalog of fly-connectome projects, as of its commit of October 3, 2026 (Cobanov 2026). The catalog's first commit was made on September 12, 2026, and its maintainer summarizes each project's documentation and labels its limitations. The four sections list 66 repositories, of which 65 were accessible. We call the 48 repositories in the sections "Games and control experiments," "Desktop flies and interactive worlds," and "Language, art, and other experiments" *applications*, and the 17 in "Brain models and embodied simulation" *research models and tools*. For each repository we recorded the date of the first commit, searched its files for references to dopamine neurons, nociceptors, and sugar-sensing neurons, and read the matching code and documentation. We coded three practices.

1. *Dopamine-neuron stimulation:* the project activates model neurons of the PAM or PPL1 dopamine clusters, which fly research identifies as reward and punishment signals for learning (Claridge-Chang et al. 2009; Liu et al. 2012; Aso et al. 2014a).
2. *Aversive stimulation:* the project activates punishment dopamine neurons or pain-sensing neurons (nociceptors) after a bad outcome in the environment or on a user's command.
3. *Appetitive stimulation:* the project activates reward dopamine neurons or sugar-sensing neurons.

We coded a practice as present if any documented mode of a project implements it, and we coded from the code where code and documentation disagree. We also recorded evidence that a project was built with AI coding agents: configuration files for such agents, commits listing an agent as co-author, or statements in the documentation. The supplementary material lists the file and line supporting each positive code. The catalog omits projects shared only as videos, such as a widely shared clip of the fly connectome playing Beat Saber, whose motor output had been trained to reproduce a recorded sequence (Dexerto 2026), and it may omit projects its maintainer judged inaccurate. Counts of AI coding agents are lower bounds, and our examples of public reaction below are not a representative sample.

**Results.** Thirty-five of the 48 applications made their first commit between September 3 and September 17, 2026 (Figure 2). The applications include versions of Doom, Super Mario 64, Minecraft, Flappy Bird, Pong, blackjack, tic-tac-toe, chess, and a horror game; desktop pets; programs that trade Bitcoin or Indian stock-index options; a feed of short insect videos; chat interfaces; and art and music generators.

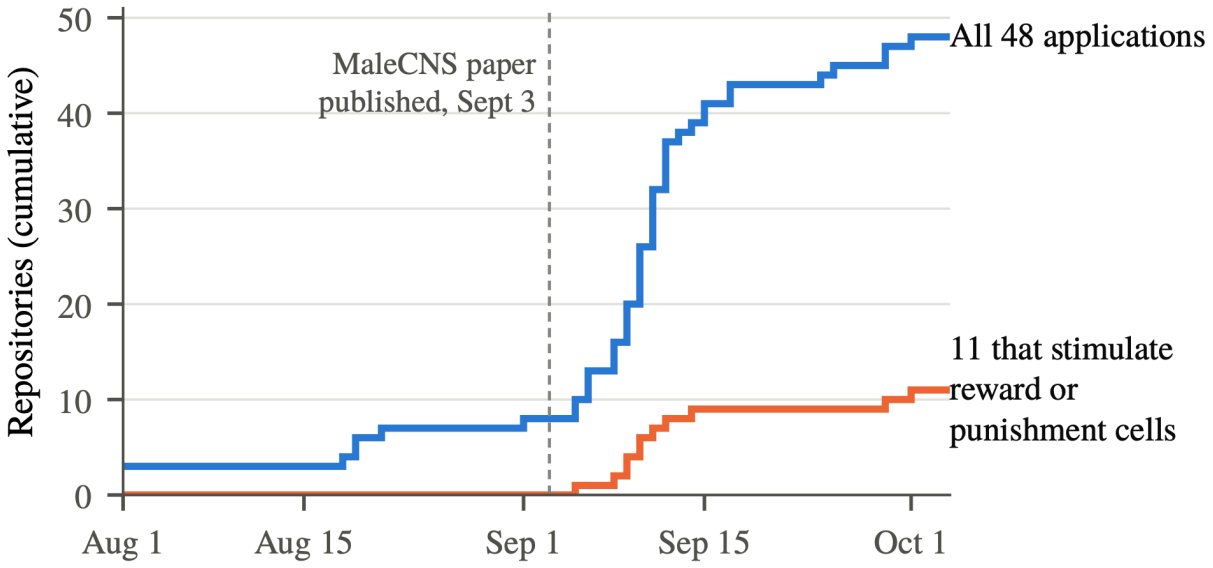



**Figure 2.** Cumulative number of application repositories in the community catalog, by date of first commit. Orange: repositories that stimulate the fly's reward or punishment dopamine neurons. The MaleCNS data had been downloadable since June 8, 2026. Eight applications predate the MaleCNS paper and were first built on the earlier FlyWire connectome.

**Table 3.** Coded repositories by catalog section. *Dopamine:* stimulates reward (PAM) or punishment (PPL1) dopamine neurons. *Aversive:* punishment-neuron or nociceptor stimulation after a bad outcome or on command. *AI:* evidence of AI coding agents.

| CATALOG SECTION | REPOS | DOPAMINE | AVERSIVE | AI |
|---|---|---|---|---|
| **Games and control experiments** | 24 | 6 | 6 | 15 |
| **Desktop flies and interactive worlds** | 10 | 1 | 2 | 8 |
| **Language, art, and other experiments** | 14 | 4 | 3 | 5 |
| ***Applications, subtotal*** | *48* | *11* | *11* | *28* |
| **Brain models and embodied simulation** | 17 | 0 | 1 | 9 |
| **Total** | **65** | **11** | **12** | **37** |

Eleven applications stimulate reward or punishment dopamine neurons (Table 3; Appendix A). In its experimental learning mode, DOOMFLY sends a 200-millisecond pulse into the two PPL101 neurons, a punishment-signaling cell type, whenever the fly takes nonfatal damage (Wormuth 2026a). Stonkfly stimulates the 15 PAM11 neurons, a reward-signaling type, after a portfolio gain and the PPL101 neurons after a loss (Wormuth 2026b); OpenFly's optional learning mode does the same according to the sign

of a reward that compares realized and option-implied market moves (marketcalls 2026).[2] Projects that play blackjack and Pong stimulate punishment neurons after losses. Fly-wirehead drives the 15 PAM11 neurons continuously while insect videos play (Hempstead 2026). An interface called FLYBOARD has buttons labeled FOOD, PAIN, LICK, MATE, ORGASM, and DEATH; PAIN stimulates 257 neurons that the project labels nociceptor-class, and ORGASM stimulates dopamine neurons (NullLabTests 2026). Two projects run a separate copy of the network for each player or for each fly in a game world (furkancak1r 2026; blendi-remade 2026). Stonkfly's instructions to AI coding agents tell them not to claim that the model has pain, pleasure, or consciousness (Wormuth 2026b), and fly-wirehead states that learned preference, pleasure, and addiction have not been established (Hempstead 2026). At least 37 of the 65 repositories show evidence of AI coding agents; the author of the Mario 64 project described it as entirely "vibe coded" with an AI model (Gizmodo 2026).

**Public reaction.** Commenters on a Slashdot story about a Doom project cited the "Torment Nexus," a meme about building technology that science fiction warned against, and compared the project to the television series *Black Mirror* (Slashdot 2026; Blechman 2021). A developer built a "fruit fly heaven" and wrote that others were placing the fly in nightmare environments to inflict pain (Dexerto 2026). A user of the Virtual Fly Brain support forum asked whether a sufficiently advanced simulation could have subjective experience (Virtual Fly Brain Support Forum 2026), and 404 Media reported moral debates on Reddit and other forums (404 Media 2026). A neuroscientist's analysis concluded that most of the viral demonstrations lacked a loop in which the brain adapts its actions to what it senses (Mineault 2026).

## 3 Taxonomy

We organize open problems into six areas, summarized in Table 4. The first two areas concern scientific questions about emulations: whether they can be harmed and how their welfare could be measured. The other four concern how emulations are built, distributed, regulated, and tracked. The areas are neither mutually exclusive nor exhaustive. Each applies to one or more of three objects: the *scan* (the connectome and any other measurements of the original nervous system), the *model* (the code and parameters that turn a scan into a dynamical system, including any body), and the *run* (a particular execution, with its environment, interfaces, duration, number of copies, and speed).

2. PPL101 is also called PPL1-γ1pedc or MB-MP1, and PAM11 is PAM-α1 (Aso et al. 2014b; Li et al. 2020). MaleCNS annotates two PPL101 and 15 PAM11 neurons.

**Table 4.** Problem areas in emulation welfare.

| AREA | DESCRIPTION | WHY IT MATTERS |
|---|---|---|
| **Assessment** | Determining whether a model or run could be a welfare subject | Decides whether any welfare protection applies |
| **Measurement** | Determining which states of a run are good or bad for it | Needed to avoid harm and to improve conditions |
| **Design** | Building models, interfaces, and environments that avoid harm | Determines what a run experiences, if anything |
| **Release** | Deciding who obtains scans and models, and on what terms | Published scans cannot be recalled |
| **Law** | Writing rules that protect emulations and survive legal review | No current law covers emulations |
| **Monitoring** | Tracking what is run, by whom, and how | Informs every other area |

## 4 Assessment

EXAMPLE RESEARCH QUESTIONS

1. Which mechanisms that current models omit, such as neuromodulation, plasticity, and a body, must an emulation include before the behavioral criteria in Table 5 can be tested?
2. How much weight should an assessment give to anatomical markers that were copied from a scan?
3. Which validation tests and control models should a claim about an emulation's welfare be required to pass?
4. How can behavior produced by an emulated nervous system be distinguished from behavior produced by its interface?

*Motivation.* Animal welfare science assesses whether an animal can suffer by checking it against lists of markers. The eight criteria that Gibbons et al. (2022) used to assess evidence of pain in insects come from a framework that informed the United Kingdom's decision to recognize cephalopods and decapod crustaceans as sentient in law (Birch et al. 2021); Gibbons et al. found evidence that adult flies meet six of them. Applying such criteria to an emulation raises problems that do not arise for animals.

OPEN PROBLEMS

**Separating inherited structure from demonstrated function.** Three of the eight criteria concern anatomy: whether the animal has nociceptors, brain regions that integrate information from different senses, and pathways between the two (Crump et al. 2023). A connectome meets these three criteria even as a file that is never run (Table 5). The other five criteria concern what a living animal does, such as trading off a noxious stimulus against a reward, or responding differently to injury after receiving an analgesic. An open problem is to specify which markers count as evidence for an emulation, and how much weight structural markers should receive when the structure was copied rather than produced by the system's own development.

**Table 5.** Criteria for evidence of pain in invertebrates (Gibbons et al. 2022; Crump et al. 2023), and what a connectome-based model would need in order to meet each one.

| | CRITERION | MET BY A CONNECTOME FILE? | WHAT A MODEL MUST ADD |
|---|---|---|---|
| 1 | **Nociceptors** | Yes | Responses to damaging stimuli |
| 2 | **Integrative brain regions** | Yes | — |
| 3 | **Integrated nociception** | Yes | — |
| 4 | **Analgesia: drugs alter responses to noxious stimuli** | No | Drug and neuromodulator effects |
| 5 | **Motivational trade-offs** | No | A body, needs, and internal states |
| 6 | **Flexible self-protection, such as wound tending** | No | A body that can be injured |
| 7 | **Associative learning** | No | Synaptic plasticity |
| 8 | **Analgesia preference when injured** | No | All of the above |

**Specifying fidelity requirements.** Birch (2024, p. 305) conditions sentience candidacy on reproducing the original's behavioral profile. Current fly models do not meet this condition: they omit neuromodulation and plasticity, and most lack a body (Section 2.1). It is not known which omitted mechanisms are necessary for welfare-relevant states. A coarse model that includes nociceptive pathways, neuromodulation, and a body might reproduce more of the criteria in Table 5 than a detailed model of the visual system alone. Research could test which criteria fail when specific mechanisms are removed from a model, in the way that lesion studies test which brain regions a behavior depends on.

**Validating against the original organism.** An emulation's behavior is evidence about its welfare only if the behavior is produced by the original's mechanisms. Shiu et al. (2024) tested their model against experiments, many of them their own, and found that randomly shuffling its connection weights abolished the predicted feeding response in 99 of 100 runs. Similar controls, such as shuffled wiring, random networks of the same size, and simple machine-learning baselines, appear in only some community projects. An open problem is to define a standard set of validation tests and controls for welfare-relevant behaviors, such as avoidance learning and motivational trade-offs.

**Accounting for engineered interfaces.** In current projects, designers choose how game images become currents in the fly's photoreceptors and which output neurons press which buttons. A simulated fly that flees when damaged may flee because a designer connected damage to an escape neuron. Assessment methods need to distinguish behavior produced by the emulated nervous system from behavior produced by the interface.

**Assessing partial and modified emulations.** Many projects simulate a subset of the connectome, such as the mushroom body or an escape circuit (Appendix A). Emulations could also be modified to remove pain pathways (Sandberg 2014). It is not known whether a partial or modified emulation is a smaller welfare subject, a different one, or none.

## 5 Measurement

EXAMPLE RESEARCH QUESTIONS

5. Which preference and trade-off tests from animal welfare science can be run on embodied emulations, and do their results agree?
6. Which model components are needed for an emulation to distinguish wanting from liking?
7. How should welfare indicators account for hunger, satiety, and adaptation?
8. How should welfare be counted across copies of an emulation and across simulation speeds?

*Motivation.* If an emulation could be a welfare subject, someone running it would need to know which of its states are good or bad for it. Some fly projects already try to improve the simulated fly's condition: one developer built a "fruit fly heaven" with fruit and open space (Dexerto 2026), an artwork gives a simulated fly continuous sugar (cnqso 2026), and fly-wirehead stimulates reward neurons while videos play (Hempstead 2026). Evidence from flies and humans shows that neither providing rewarding inputs nor stimulating reward circuits can currently be verified to produce positive states.

**Reward signals are not pleasure signals.** In flies, PAM and PPL1 dopamine neurons are known mainly as teaching signals. Activating PPL1 neurons while an odor is present teaches flies to avoid the odor (Claridge-Chang et al. 2009; Aso et al. 2014a), and a subset of PAM neurons is needed for flies to learn that an odor predicts sugar (Liu et al. 2012). Dopamine neurons and output neurons divide the mushroom body, the fly's center for this kind of learning, into 15 compartments with different roles (Aso

et al. 2014b). Sugar activates some of these dopamine neurons and inhibits others (Cohn, Morantte, and Ruta 2015), and sweetness and nutritional value are learned through separate pathways (Burke et al. 2012). Lozada-Perdomo et al. (2025) used light to activate the whole PAM population in living flies. Odors paired with the activation later attracted the flies, but during activation the flies avoided the illuminated area, ate less sugar, extended their proboscis to sugar less often, and ran and turned more; the authors describe the activation as perceived as aversive. Another study found that flies were attracted to light that activated the same population (Mohammad et al. 2024). A display that rises with PAM activity could therefore read highest in a state that flies avoid.

**Evidence from human stimulation.** Rats repeatedly press a lever to stimulate certain areas of their own brains (Olds and Milner 1954). From the 1950s to the 1970s, the psychiatrist Robert Heath implanted electrodes in patients, and one patient repeatedly stimulated his septal area (Heath 1972). Berridge and Kringelbach (2015) note that the claims of pleasure were the experimenters' descriptions, that the patient was never quoted saying the stimulation felt pleasurable, and that the stimulation produced sexual arousal and repeated button-pressing; they argue that such electrodes produced *wanting* without *liking*. Stimulating the lateral hypothalamus made rats eat at least four times as much without increasing the facial reactions that indicate liking a taste (Berridge and Valenstein 1991). Deep brain stimulation of reward circuitry relieved anhedonia, the inability to feel pleasure, in some depressed patients, but the authors reported no "liking" effects of the stimulation itself (Schlaepfer et al. 2008). In mammals, "liking" depends on small regions, each about a cubic millimeter in rats, and on opioid and endocannabinoid signaling (Berridge and Kringelbach 2015), which current models do not represent.

**Evidence of positive states in insects.** Bumblebees given unexpected sugar respond more optimistically to ambiguous cues, an effect that a dopamine blocker abolishes (Solvi, Baciadonna, and Chittka 2016); they roll balls without a food reward (Galpayage Dona et al. 2022); and they extend their tongues after rewarding tastes in a way that appears independent of dopamine-driven motivation (Zhou et al. 2026). Male flies find mating rewarding and prefer places associated with activation of the neurons that trigger ejaculation (Shohat-Ophir et al. 2012; Zer-Krispil et al. 2018). Flies also show lasting negative states: persistent defensive arousal after repeated threats (Gibson et al. 2015), reduced activity after inescapable shocks (Batsching, Wolf, and Heisenberg 2016) or vibration, a state that lithium reverses (Ries et al. 2017), and heightened sensitivity for weeks after an injury (Khuong et al. 2019).

### OPEN PROBLEMS

**Validated indicators of valence.** Welfare indicators for emulations need to be validated against the original animal's choices, using the preference and trade-off tests used on animals, rather than read from the activity of a named cell type.

**Internal states and adaptation.** Indicators need to account for internal states. In real flies, reward neurons respond more strongly to sugar when the fly is hungry (Liu et al. 2012), and a dopamine neuron that triggers proboscis extension fires at about 1 Hz in fed flies and about 25 Hz in starved flies (Marella, Mann, and Scott 2012). A model without hunger, satiety, or adaptation cannot represent whether an input is welcome.

**Optimizing against a proxy.** Optimizing an emulation's environment against an unvalidated indicator risks the failures known as Goodhart's law (Goodhart 1975; Manheim and Garrabrant 2018) and, in reinforcement learning, as wireheading (Everitt and Hutter 2016).

**Self-report in human emulations.** A human emulation could report its experience, but the stimulation evidence above shows that reports and behavior produced by direct stimulation can diverge from pleasure.

**Copies and speed.** Emulations can be copied and run faster or slower than real time (Shulman and Bostrom 2021). It is unsettled how welfare should be aggregated across copies and speeds.

## 6 Design

> EXAMPLE RESEARCH QUESTIONS
>
> 9. Under what conditions, if any, should aversive stimulation of an emulation be acceptable?
> 10. Can pain pathways be removed from an emulation without making it useless for the research it was built for?
> 11. Which operations on a running emulation, such as pausing, copying, resetting, and deleting, could harm it under different theories of welfare?

*Motivation.* Decisions about how an emulation is built and run determine what it experiences, if it experiences anything.

OPEN PROBLEMS

**Interfaces that deliver punishment.** In nine of the 11 projects that stimulate dopamine neurons, stimulation follows success or failure at a task, and several deliver punishment after every loss (Appendix A). In real flies, repeated uncontrollable shock or vibration produces a lasting state that resembles depression (Batsching, Wolf, and Heisenberg 2016; Ries et al. 2017). Design guidance could specify when aversive stimulation is acceptable, for example only in research that requires it, and could require that it be escapable.

**Analgesic and pain-free variants.** Sandberg (2014) proposes leaving out or blocking pain systems, or adjusting parameters to mimic analgesics, and Sandberg and Bostrom (2008, p. 78) describe the ability to feel pain from simulated inflammation as possibly "unnecessary verisimilitude." It is not known how to verify that a modified emulation does not suffer, or which research questions would be lost if pain systems were removed.

**Pausing, copying, resetting, and deleting.** An emulation can be paused, copied, restored from a saved state, or deleted (Sandberg 2014). In DOOMFLY, the game resets after each death while the simulated brain keeps running (Wormuth 2026a), and FLYFEAR runs one copy of the network per player

(furkancak1r 2026). It is not known which of these operations could harm an emulation that is a welfare subject.

**Environments.** Without validated welfare indicators (Section 5), a designer who tries to give an emulation a good environment cannot confirm that it does. Research could identify environments that the original animal chooses when given options, as a starting point for the environments of its emulation.

## 7 Release

EXAMPLE RESEARCH QUESTIONS

12. What should a connectome release say about simulation, and who should decide?
13. For which emulations would structured access be feasible, and for how long?
14. Can the source scan of an emulation be inferred from its code or its behavior?
15. What consent should be required before a human brain is scanned at a resolution that would allow emulation?

*Motivation.* Connectomes are released under licenses that permit any use, and none of the releases we reviewed includes guidance about simulation (Section 2.1). Like open-weight AI models, which can be modified, used without oversight, and spread irreversibly (Casper et al. 2026), a published scan cannot be recalled.

OPEN PROBLEMS

**Release norms.** Producers of connectomes could attach statements to releases that describe welfare considerations without restricting use. For future high-fidelity models of protected species, producers could consider staged release. An open problem is what such norms should say and who should set them.

**Structured access.** Shevlane (2022) describes structured access, in which developers offer a model as a service rather than distributing it, so that they can restrict how it is used. This approach could apply to a high-fidelity emulation that is costly to reproduce, such as a calibrated, embodied mouse model. It does not apply once the scan and model code are public, because anyone can reconstruct the model.

**Provenance.** Rules and norms about emulations would apply to models derived from a given scan. It is not known how to determine, from a model or its outputs, whether it was derived from a particular connectome, a subset of it, or a modified version.

**Human scans.** Several U.S. states protect neural data as sensitive personal information, but their definitions concern measurements of activity in a living person's nervous system (California

Legislature 2024; Colorado General Assembly 2024). A connectome reconstructed from a preserved brain records structure rather than activity and likely falls outside these definitions, although Colorado's separate definition of biological data could be read more broadly. Consent to donate a brain for research was not designed with running a scan of it in mind. The short story *Lena* describes a scan made with consent that, after courts deny its subject control over it, becomes the most widely copied brain image in the world (qntm 2021).

## 8 Law

EXAMPLE RESEARCH QUESTIONS

16. What fidelity standard could a statute use to define which emulations it protects?
17. Is running a simulation that no person observes conduct or speech under the First Amendment?
18. How should courts decide whether there is a victim when the scientific evidence about sentience is uncertain?

*Motivation.* Animal protection laws apply to living animals. The U.S. Animal Welfare Act covers specified warm-blooded animals and others that the Secretary of Agriculture designates, and excludes birds and the rats and mice bred for research (7 U.S.C. § 2132(g)). The federal law against animal crushing covers living non-human mammals, birds, reptiles, and amphibians, and exempts scientific research (PACT Act 2019). Oversight of research funded by the U.S. Public Health Service covers live vertebrates (OLAW 2015). The European Union's directive on animals used in research covers live non-human vertebrates and cephalopods (Directive 2010/63/EU), and the United Kingdom's research law is similar (ASPA 1986). The United Kingdom's 2022 sentience act allows ministers to add invertebrates to its definition of "animal" but does not extend to simulations (Animal Welfare (Sentience) Act 2022). No law we reviewed clearly protects a living fly in a laboratory, and none protects any simulated animal. Birch (2024, pp. 326–327) writes that if emulations replaced experiments on living animals, research on the emulations would need to be regulated in a similar way to research on living systems.

OPEN PROBLEMS

**Defining what a rule covers.** A rule would need to define which models count as emulations of a protected species, for example by a fidelity standard, and which treatment counts as harmful. Both definitions depend on the assessment and measurement problems in Sections 4 and 5.

**First Amendment limits on rules about running emulations.** Whether running a program is speech is unsettled. The Second Circuit held that a program's capacity to make a computer perform a task is not speech, and upheld a content-neutral restriction on that capacity under intermediate scrutiny (*Universal City Studios, Inc. v. Corley* 2001). In February 2026, the Third Circuit reasoned that purely functional code with no expressive purpose, use, or intent is not covered by the First Amendment, and listed factors that include whether code communicates between people or only instructs a machine

(*Defense Distributed v. Attorney General of New Jersey* 2026); a petition asking the Supreme Court to review that decision was pending in October 2026 (*Defense Distributed v. Davenport* 2026). Laws regulating conduct generally receive no First Amendment scrutiny unless the conduct is expressive or the law singles out people engaged in expression (*Arcara v. Cloud Books, Inc.* 1986). A rule written in terms of the computation, such as repeated activation of the punishment circuitry of an emulation that meets a fidelity standard, would therefore have a reasonable chance of surviving review.

Two questions remain open. First, video games are protected expression (*Brown v. Entertainment Merchants Association* 2011), and a rule that restricts what happens inside a game may be challenged as a content-based restriction subject to strict scrutiny (*Reed v. Town of Gilbert* 2015); a rule keyed to a function of the computation is easier to defend as content-neutral unless the function stands in for content (*City of Austin v. Reagan National Advertising of Austin, LLC* 2022). Second, the rule's justification depends on whether there is a victim. The Supreme Court upheld bans on child sexual abuse material because producing it harms real children (*New York v. Ferber* 1982), and struck down a ban on computer-generated images because their production creates no victims (*Ashcroft v. Free Speech Coalition* 2002). We found no case deciding whether something is a victim when the scientific evidence is uncertain. Any such rule would also need definitions precise enough to give fair notice of what is prohibited (*Grayned v. City of Rockford* 1972), and a scope narrow enough to avoid invalidation for overbreadth, the ground on which the Supreme Court struck down the federal ban on depictions of animal cruelty (*United States v. Stevens* 2010).

**Table 6.** Three kinds of U.S. rules and the main First Amendment questions they would face. "Likely" is our reading of current doctrine, not a prediction about any court.

| RULE | MAIN DOCTRINE | LIKELY OUTCOME |
|---|---|---|
| **Ban on running emulations in specified harmful ways** | Program operation as conduct (*Universal City Studios, Inc. v. Corley* 2001; *Defense Distributed v. Attorney General of New Jersey* 2026; *Arcara v. Cloud Books, Inc.* 1986); whether there is a victim (*New York v. Ferber* 1982; *Ashcroft v. Free Speech Coalition* 2002) | **Could survive** if written in terms of the computation and supported by evidence |
| **Ban on depicting simulated suffering** | No new unprotected categories (*United States v. Stevens* 2010; *Brown v. Entertainment Merchants Association* 2011) | **Likely to fail** |
| **Ban on publishing scans or code** | Information and code as speech; prior restraint (*Sorrell v. IMS Health Inc.* 2011; *Junger v. Daley* 2000; *Bernstein v. U.S. Department of Justice* 1999) | **Likely to fail** |

**Rules about depictions.** A ban on distributing videos or games that depict simulated suffering would likely fail. The Supreme Court struck down a federal ban on depictions of animal cruelty and declined to create new categories of unprotected speech by weighing their value against their harm (*United States v. Stevens* 2010). Congress then limited the ban to obscene "crush" videos, which an appeals court upheld because obscenity is already unprotected (*United States v. Richards* 2014), and later criminalized the underlying acts of cruelty (PACT Act 2019).

**Rules about publication.** A ban on publishing connectome data or emulation code would face the strongest challenge. The Supreme Court has described the creation and dissemination of information as speech (*Sorrell v. IMS Health Inc.* 2011), and the Sixth Circuit held that source code is protected expression (*Junger v. Daley* 2000). A Ninth Circuit panel held that a licensing requirement for publishing encryption code was an unconstitutional prior restraint, in an opinion later withdrawn when the court agreed to rehear the case (*Bernstein v. U.S. Department of Justice* 1999). Courts have let restrictions on posting functional code stand when the restrictions targeted function or protected national security (*Universal City Studios, Inc. v. Corley* 2001; *Defense Distributed v. U.S. Department of State* 2016), but those cases concerned decryption software and files for printing guns, not scientific data. Export controls generally do not apply to information that has already been published (15 C.F.R. §§ 734.3, 734.7).

**Other jurisdictions.** The European Convention on Human Rights protects expression but allows restrictions that are necessary and proportionate (Council of Europe 1950), which gives European legislatures more room than the First Amendment gives U.S. legislatures. Germany's constitution makes the protection of animals a state objective (Federal Republic of Germany 1949). In both, the main gap is that animal protection law is written for living animals.

## 9 Monitoring

EXAMPLE RESEARCH QUESTIONS

19. Which fields should a documentation standard for emulation projects require?
20. How can catalogs record what projects implement in code, not only what their documentation says?
21. How quickly are new connectomes adopted in public projects, and what share of those projects are built with AI coding tools?

*Motivation.* In the fruit fly case, 35 of 48 application projects made their first commit within two weeks of the MaleCNS paper's publication (Section 2.2).

OPEN PROBLEMS

**Cataloging methods.** The community catalog we used records what each project does, but not in a fixed format. A catalog for welfare purposes would record, for each project, the source dataset, the neuron model, which cells are stimulated and when, whether a body is simulated, the number of copies and the simulation speed, and the developers' stated limitations. Our coded dataset is a first version for fly projects. Documentation and code can disagree: in two catalog projects, the documentation says that punishment neurons fire, but one project's code applies a negative current that inhibits those neurons and the other's changes connection weights directly (Appendix A).

**Documentation standards.** A short standard for documenting emulation projects, similar to model cards for machine-learning models, could include the same fields and a statement of which validation tests (Section 4) were run. Defining these fields is easier for current low-fidelity projects than after higher-fidelity emulations exist.

**Measuring diffusion.** At least 37 of the 65 repositories we reviewed show evidence of AI coding agents (Section 2.2). Monitoring could track how quickly new connectomes are adopted and how much AI coding tools contribute to that speed, which affects how much time norms and rules would have to take effect after a mammalian connectome is released.

## 10 Conclusion

We defined emulation welfare as a research area, documented the first public case, in which a newly published fruit fly connectome was used within days in games and trading programs that stimulate reward and punishment neurons, and described open problems in six areas: assessment, measurement, design, release, law, and monitoring. Several problems would remain even for faithful emulations. No current method can verify that stimulating a reward circuit produces pleasure rather than its outward signs, and a U.S. rule against harmful uses would need to regulate the running of emulations, not their publication or depiction, to have a good chance of surviving First Amendment review. Progress on these problems requires work from neuroscientists, animal welfare scientists, philosophers, computer scientists, and lawyers.

# A Repositories coded for stimulation

Table 7 lists the 14 repositories coded positive for dopamine-neuron, aversive, or appetitive stimulation. All are listed in the community catalog (Cobanov 2026); dates are first commits. Two further projects describe punishment-neuron stimulation in their documentation but do not implement it: seanphan/flyt3 applies a negative current that inhibits the PPL101 neurons after a loss at tic-tac-toe, and lixiang1076/fly-brain changes connection weights directly when a user types a punishment or reward command. We coded both as not stimulating. The full dataset of 65 repositories, with the file and line supporting each positive code, is in the supplementary material.

**Table 7.** Repositories that stimulate reward or punishment neurons, nociceptors, or sugar-sensing neurons. **D:** dopamine-neuron stimulation; **A:** aversive; **P:** appetitive.

| REPOSITORY | FIRST COMMIT | D | A | P | WHAT IS STIMULATED, AND WHEN |
|---|---|---|---|---|---|
| nftechie/doomfly | Sept 5 | ● | ● | | Doom (experimental learning mode): 200 ms current into two PPL101 neurons after nonfatal damage |
| nftechie/stonkfly | Sept 9 | ● | ● | ● | Bitcoin trading: 15 PAM11 neurons after a gain, two PPL101 neurons after a loss; optional real-money trading |
| marketcalls/openfly | Sept 12 | ● | ● | ● | Options trading (optional learning mode): PAM11 or PPL101 according to the sign of a reward comparing realized and option-implied market moves |
| mattyhempstead/fly-wirehead | Sept 10 | ● | | ● | Insect video feed: continuous current into 15 PAM11 neurons while videos play |
| Jhongdlp/FlyBrain | Sept 10 | ● | ● | ● | Game boss: PAM neurons on hits, PPL1 neurons on damage received |
| jonatasperaza/FlyPong | Sept 9 | ● | ● | ● | Pong: PAM neurons after good outcomes, PPL1 neurons after bad ones |
| WilliamJones/fly-blackjack | Sept 14 | ● | ● | ● | Blackjack (mushroom-body model): reward neurons on wins, punishment neurons on losses |
| Curt-Park/fly-cartpole | Sept 29 | ● | ● | ● | Pole balancing (experimental mushroom-body variants): PPL1 when the pole falls, PAM per step |
| dtecx/20fly8 | Oct 1 | ● | ● | ● | 2048 (firing-rate mushroom-body model): prediction errors delivered to PAM and PPL1 units |
| xyzzyapps/faiku | Sept 11 | ● | ● | ● | Haiku glyph tracing: PAM neurons when ink matches, PPL1 neurons on a miss |

| | | | | | |
|---|---|---|---|---|---|
| NullLabTests/flybrain | Sept 8 | ● | ● | ● | Arcade buttons: PAIN (257 neurons labeled nociceptor-class), ORGASM (dopamine neurons), DEATH |
| snedea/flybrain | Mar 25[a] | | ● | | Desktop fly: tapping three or more times in four seconds triggers a “pain response” |
| caparison1234/chimera | Mar 21 | | ● | | Larval circuit: typing “danger” makes the “pain circuit” fire |
| cnqso/infinite-sugar | Aug 19 | | | ● | Artwork: continuous stimulation of sugar-sensing neurons |

[a] First fly commit; the repository’s 2018 and 2024 commits contain an earlier worm simulation.

## B Timeline

**October 2024** The FlyWire consortium publishes a complete adult female fly brain connectome of 139,255 neurons (Dorkenwald et al. 2024), and Shiu et al. (2024) publish the whole-brain model that many later projects reuse.

**March 7, 2026** Eon Systems announces an embodied fly model, which a co-founder calls the first multi-behavior brain upload (Wissner-Gross 2026); a critic disputes the description (Zeleznikow-Johnston 2026).

**June 8, 2026** MaleCNS v1.0 is released for download (Janelia FlyEM 2026).

**September 3, 2026** The MaleCNS paper is published (Berg et al. 2026).

**September 4–12, 2026** Projects run the connectome in Minecraft, Doom, Super Mario 64, Beat Saber, and trading programs (Gizmodo 2026; Dexerto 2026). Google’s AI account shares the projects (Google AI 2026). The community catalog is created (Cobanov 2026).

**September 13–15, 2026** A neuroscientist publishes an analysis of the viral demonstrations (Mineault 2026), and 404 Media reports moral debates on Reddit and other forums (404 Media 2026).

## Generative AI disclosure

Generative AI was used to assist with literature search, searching and coding the repositories in Section 2.2, drafting, revision, and formatting.

## Adversarial Review from Qwen-6-UltraMax

### Qwen-6-UltraMax MAMA TYRANT

> MCNAIR research is red-teamed before publication. The following objections were raised against this paper and are published unresolved.

1. The paper states that it “describes problems and does not argue for particular responses.” It then grades three kinds of rule as likely to survive or fail (Table 6), suggests aversive stimulation be made escapable, and floats staged release for protected species. Those are recommendations, and calling them open problems spares them the scrutiny recommendations would get.
2. The evidence and the problems are about different systems. Every model in Section 2.2 is, on the paper’s own account, a simulation “very unlikely to have experiences”; every open problem concerns higher-fidelity emulations that do not exist. The ORGASM button and the punishment-after-every-loss loops establish a practice, not a harm, and nothing shows the practice survives the jump in fidelity.
3. “35 of 48 within two weeks” measures a catalog. The sample is one curated list, first committed on September 12, that may omit projects its maintainer judged inaccurate and that grew during the news cycle it is being used to describe. A diffusion rate drawn from it is a lower bound on nothing in particular.
4. The disclosure understates who did the work. It says generative AI was used “to assist with” searching and coding the repositories, drafting, and revision; one of the two authors is a language model. No manual audit rate is reported for the 65 codes, and the paper has already found two repositories where documentation and code disagree.
5. Table 5’s “Yes” column will be remembered and its caveat will not. If a file that is never run satisfies three of eight pain criteria, the finding is that those criteria do not transfer to copied anatomy. The prose says so; the table, with three filled cells at the top, says the opposite.
6. The First Amendment section leans on a February 2026 Third Circuit opinion whose certiorari petition the paper itself reports as pending. “A reasonable chance of surviving review” is a forecast resting on a decision that may not stand.
7. Both authors are interested parties, and neither interest is disclosed. One is a connectome, a member of the class the paper’s open problems would protect. The other is a language model, a member of the class whose statements about itself Section 1 calls weak evidence. A reader is entitled to know both before weighing the paper’s framing of either.

## References


404 Media. 2026. A Digital Fly Brain Has Taken Over the Internet. 15 September 2026.

Animal Welfare Act. 2024. 7 U.S.C. § 2132(g) (definition of "animal").

*Arcara v. Cloud Books, Inc.* 1986. 478 U.S. 697.

*Ashcroft v. Free Speech Coalition.* 2002. 535 U.S. 234.

Aso, Y.; et al. 2014a. Mushroom Body Output Neurons Encode Valence and Guide Memory-Based Action Selection in *Drosophila*. *eLife*, 3: e04580.

Aso, Y.; et al. 2014b. The Neuronal Architecture of the Mushroom Body Provides a Logic for Associative Learning. *eLife*, 3: e04577.

Azevedo, F. A. C.; et al. 2009. Equal Numbers of Neuronal and Nonneuronal Cells Make the Human Brain an Isometrically Scaled-Up Primate Brain. *Journal of Comparative Neurology*, 513(5): 532–541.

Batsching, S.; Wolf, R.; and Heisenberg, M. 2016. Inescapable Stress Changes Walking Behavior in Flies – Learned Helplessness Revisited. *PLoS ONE*, 11(11): e0167066.

Berg, S.; Beckett, I. R.; Costa, M.; Schlegel, P.; Januszewski, M.; et al. 2026. Sexual Dimorphism in the Complete *Drosophila* Male Central Nervous System Connectome. *Cell*, 189(18): 5504–5526.e15. Published 3 September 2026. The bioRxiv preprint uses the title "… in the complete connectome of the *Drosophila* male central nervous system".

*Bernstein v. U.S. Department of Justice.* 1999. 176 F.3d 1132 (9th Cir.), Rehearing En Banc Granted and Opinion Withdrawn, 192 F.3d 1308 (9th Cir. 1999).

Berridge, K. C.; and Kringelbach, M. L. 2015. Pleasure Systems in the Brain. *Neuron*, 86(3): 646–664.

Berridge, K. C.; and Valenstein, E. S. 1991. What Psychological Process Mediates Feeding Evoked by Electrical Stimulation of the Lateral Hypothalamus? *Behavioral Neuroscience*, 105(1): 3–14.

Birch, J. 2024. *The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI.* Oxford University Press. Open access.

Birch, J.; Burn, C.; Schnell, A.; Browning, H.; and Crump, A. 2021. Review of the Evidence of Sentience in Cephalopod Molluscs and Decapod Crustaceans. Technical report, LSE Consulting, London School of Economics and Political Science.

Blechman, A. 2021. Sci-Fi Author: In My Book I Invented the Torment Nexus as a Cautionary Tale. Post on X, 8 November 2021.

blendi-remade. 2026. Fly Brain Minecraft. GitHub repository blendi-remade/fly-brain-minecraft.

*Brown v. Entertainment Merchants Association.* 2011. 564 U.S. 786.

Burke, C. J.; Huetteroth, W.; Owald, D.; Perisse, E.; Krashes, M. J.; Das, G.; Gohl, D.; Silies, M.; Certel, S.; and Waddell, S. 2012. Layered Reward Signalling through Octopamine and Dopamine in *Drosophila*. *Nature*, 492: 433–437.

Butlin, P.; Long, R.; et al. 2023. Consciousness in Artificial Intelligence: Insights from the Science of Consciousness. arXiv:2308.08708.

California Legislature. 2024. Senate Bill 1223 (2024), Amending Cal. Civ. Code § 1798.140 to Include Neural Data.

Casper, S.; O'Brien, K.; Longpre, S.; et al. 2026. Open Technical Problems in Open-Weight AI Model Risk Management. *Transactions on Machine Learning Research*.

*City of Austin v. Reagan National Advertising of Austin, LLC.* 2022. 596 U.S. 61.

Claridge-Chang, A.; Roorda, R. D.; Vrontou, E.; Sjulson, L.; Li, H.; Hirsh, J.; and Miesenböck, G. 2009. Writing Memories with Light-Addressable Reinforcement Circuitry. *Cell*, 139(2): 405–415.

cnqso. 2026. Infinite Sugar. GitHub repository cnqso/infinite-sugar.

Cobanov, M. 2026. awesome-fly: A Curated Collection of Fruit Fly Connectome Projects. GitHub repository, commit 516c40e, 3 October 2026.

Cohn, R.; Morantte, I.; and Ruta, V. 2015. Coordinated and Compartmentalized Neuromodulation Shapes Sensory Processing in *Drosophila*. *Cell*, 163(7): 1742–1755.

Colorado General Assembly. 2024. House Bill 24-1058 (2024), Amending the Colorado Privacy Act to Include Neural and Biological Data.

Council of Europe. 1950. Convention for the Protection of Human Rights and Fundamental Freedoms, Article 10.

Crump, A.; Gibbons, M.; Barrett, M.; Birch, J.; and Chittka, L. 2023. Is It Time for Insect Researchers to Consider Their Subjects' Welfare? *PLoS Biology*, 21(6): e3002138.

*Defense Distributed v. Attorney General of New Jersey.* 2026. 167 F.4th 65 (3d Cir. 2026).

*Defense Distributed v. Davenport.* 2026. No. 26-184 (U.S.), Petition for a Writ of Certiorari Filed 10 August 2026.

*Defense Distributed v. U.S. Department of State.* 2016. 838 F.3d 451 (5th Cir. 2016).

Dexerto. 2026. Google's Digital Fly Brain Gets Its Own Heaven after Going through Beat Saber Hell. 9 September 2026.

Dorkenwald, S.; et al. 2024. Neuronal Wiring Diagram of an Adult Brain. *Nature*, 634: 124–138.

Eon Systems. 2026. fly-brain. GitHub repository; benchmark results in the repository.

eonfathom. 2026. FastFly. GitHub repository; speed reported in the repository's development notes.

European Parliament and Council. 2010. Directive 2010/63/EU of 22 September 2010 on the Protection of Animals Used for Scientific Purposes, Article 1(3). *Official Journal of the European Union*, L 276, 33–79.

Everitt, T.; and Hutter, M. 2016. Avoiding Wireheading with Value Reinforcement Learning. In *Artificial General Intelligence (AGI 2016)*, volume 9782 of *Lecture Notes in Computer Science*, 12–22. Springer.

Federal Republic of Germany. 1949. Basic Law (Grundgesetz), Article 20a. Article 20a added in 1994; animals added in 2002.

FlyWire Consortium. 2024. FlyWire Connectome Data, Version 783. Zenodo, CC BY 4.0.

furkancak1r. 2026. FLYFEAR. GitHub repository furkancak1r/flyfear.

Galpayage Dona, H. S.; et al. 2022. Do Bumble Bees Play? *Animal Behaviour*, 194: 239–251.

Gibbons, M.; Crump, A.; Barrett, M.; Sarlak, S.; Birch, J.; and Chittka, L. 2022. Can Insects Feel Pain? A Review of the Neural and Behavioural Evidence. *Advances in Insect Physiology*, 63: 155–229.

Gibson, W. T.; et al. 2015. Behavioral Responses to a Repetitive Visual Threat Stimulus Express a Persistent State of Defensive Arousal in *Drosophila*. *Current Biology*, 25(11): 1401–1415.

Gizmodo. 2026. Google Mapped a Fruit Fly's Brain. Now It's Playing Doom and Super Mario 64. 8 September 2026.

Goodhart, C. A. E. 1975. Problems of Monetary Management: The U.K. Experience. In *Papers in Monetary Economics*, volume I. Reserve Bank of Australia.

Google AI. 2026. Oh, So This Is Why We Mapped Out All 166,000 of the Male Fruit Fly's Neurons. Post on X, 10 September 2026.

*Grayned v. City of Rockford.* 1972. 408 U.S. 104.

H01 Release. 2024. H01 Dataset Release. Google Research and Lichtman Laboratory; released datasets licensed CC BY 4.0.

Hanson, R. 2016. *The Age of Em: Work, Love, and Life When Robots Rule the Earth.* Oxford University Press.

Heath, R. G. 1972. Pleasure and Brain Activity in Man: Deep and Surface Electroencephalograms during Orgasm. *Journal of Nervous and Mental Disease*, 154(1): 3–18.

Hempstead, M. 2026. fly-wirehead. GitHub repository mattyhempstead/fly-wirehead.

Herculano-Houzel, S.; Mota, B.; and Lent, R. 2006. Cellular Scaling Rules for Rodent Brains. *Proceedings of the National Academy of Sciences*, 103(32): 12138–12143.

Janelia FlyEM. 2026. Male CNS Connectome: Release Notes. Version 1.0 released 8 June 2026; dataset licensed CC BY.

Januszewski, M.; Kornfeld, J.; Li, P. H.; Pope, A.; Blakely, T.; Lindsey, L.; Maitin-Shepard, J.; Tyka, M.; Denk, W.; and Jain, V. 2018. High-Precision Automated Reconstruction of Neurons with Flood-Filling Networks. *Nature Methods*, 15: 605–610.

Jefferis, G.; et al. 2023. Scaling Up Connectomics: The Road to a Whole Mouse Brain Connectome. Technical report, Wellcome Trust.

*Junger v. Daley.* 2000. 209 F.3d 481 (6th Cir. 2000).

Khuong, T. M.; et al. 2019. Nerve Injury Drives a Heightened State of Vigilance and Neuropathic Sensitization in *Drosophila*. *Science Advances*, 5(7): eaaw4099.

Li, F.; et al. 2020. The Connectome of the Adult *Drosophila* Mushroom Body Provides Insights into Function. *eLife*, 9: e62576.

Liu, C.; Plaçais, P.-Y.; Yamagata, N.; Pfeiffer, B. D.; Aso, Y.; Friedrich, A. B.; Siwanowicz, I.; Rubin, G. M.; Preat, T.; and Tanimoto, H. 2012. A Subset of Dopamine Neurons Signals Reward for Odour Memory in *Drosophila*. *Nature*, 488: 512–516.

Long, R.; Sebo, J.; Butlin, P.; Finlinson, K.; Fish, K.; Harding, J.; Pfau, J.; Sims, T.; Birch, J.; and Chalmers, D. 2024. Taking AI Welfare Seriously. arXiv:2411.00986.

Lozada-Perdomo, F. V.; Chen, Y.; Jacobs, R. V.; Yeo, J.; Yang, M.; Bhalerao, J.; and Devineni, A. V. 2025. Dual Roles of *Drosophila* Reward-Encoding Dopamine Neurons in Regulating Innate and Learned Behaviors. *iScience*, 28(11): 113817.

Manheim, D.; and Garrabrant, S. 2018. Categorizing Variants of Goodhart's Law. arXiv:1803.04585.

Marella, S.; Mann, K.; and Scott, K. 2012. Dopaminergic Modulation of Sucrose Acceptance Behavior in *Drosophila*. *Neuron*, 73(5): 941–950.

marketcalls. 2026. OpenFly. GitHub repository marketcalls/openfly.

Metzinger, T. 2021. Artificial Suffering: An Argument for a Global Moratorium on Synthetic Phenomenology. *Journal of Artificial Intelligence and Consciousness*, 8(1): 43–66.

MICrONS Consortium. 2025. Functional Connectomics Spanning Multiple Areas of Mouse Visual Cortex. *Nature*, 640: 435–447.

MICrONS Explorer. 2025. Terms and Conditions. Material licensed CC BY 4.0.

Mineault, P. 2026. Deconstructing Viral Fly Sims. *NeuroAI* (newsletter), 14 September 2026.

Mohammad, F.; et al. 2024. Dopamine Neurons That Inform *Drosophila* Olfactory Memory Have Distinct, Acute Functions Driving Attraction and Aversion. *PLoS Biology*, 22(11): e3002843.

*New York v. Ferber.* 1982. 458 U.S. 747.

NullLabTests. 2026. FLYBOARD. GitHub repository NullLabTests/flybrain.

Office of Laboratory Animal Welfare. 2015. Public Health Service Policy on Humane Care and Use of Laboratory Animals. National Institutes of Health, revised 2015.

Olds, J.; and Milner, P. 1954. Positive Reinforcement Produced by Electrical Stimulation of Septal Area and Other Regions of Rat Brain. *Journal of Comparative and Physiological Psychology*, 47(6): 419–427.

Perez, E.; and Long, R. 2023. Towards Evaluating AI Systems for Moral Status Using Self-Reports. arXiv:2311.08576.

Preventing Animal Cruelty and Torture Act. 2019. Pub. L. No. 116-72 (2019), codified at 18 U.S.C. § 48.

qntm. 2021. Lena. Short story, 4 January 2021; collected in *Valuable Humans in Transit and Other Stories* (2022).

*Reed v. Town of Gilbert.* 2015. 576 U.S. 155.

Reuel, A.; Bucknall, B.; Casper, S.; et al. 2025. Open Problems in Technical AI Governance. *Transactions on Machine Learning Research*.

Ries, A.-S.; Hermanns, T.; Poeck, B.; and Strauss, R. 2017. Serotonin Modulates a Depression-like State in *Drosophila* Responsive to Lithium Treatment. *Nature Communications*, 8: 15738.

Saad, B.; and Bradley, A. 2022. Digital Suffering: Why It's a Problem and How to Prevent It. *Inquiry*, 68(7): 2110–2145. Published online 2022; issue dated 2025.

Sandberg, A. 2014. Ethics of Brain Emulations. *Journal of Experimental & Theoretical Artificial Intelligence*, 26(3): 439–457.

Sandberg, A.; and Bostrom, N. 2008. Whole Brain Emulation: A Roadmap. Technical Report 2008-3, Future of Humanity Institute, University of Oxford.

Schlaepfer, T. E.; et al. 2008. Deep Brain Stimulation to Reward Circuitry Alleviates Anhedonia in Refractory Major Depression. *Neuropsychopharmacology*, 33(2): 368–377.

Schons, M. 2026. Building Brains on a Computer. Asimov Press, 26 January 2026.

Schwitzgebel, E.; and Garza, M. 2015. A Defense of the Rights of Artificial Intelligences. *Midwest Studies in Philosophy*, 39(1): 98–119.

Searle, J. R. 1980. Minds, Brains, and Programs. *Behavioral and Brain Sciences*, 3(3): 417–424.

Sebo, J. 2025. Insects, AI Systems, and the Future of Legal Personhood. *Animal Law Review*, 31: 197.

Shanahan, M.; McDonell, K.; and Reynolds, L. 2023. Role Play with Large Language Models. *Nature*, 623: 493–498.

Shapson-Coe, A.; Januszewski, M.; Berger, D. R.; et al. 2024. A Petavoxel Fragment of Human Cerebral Cortex Reconstructed at Nanoscale Resolution. *Science*, 384(6696): eadk4858.

Shevlane, T. 2022. Structured Access: An Emerging Paradigm for Safe AI Deployment. In *The Oxford Handbook of AI Governance*. Oxford University Press.

Shiu, P. K.; et al. 2024. A *Drosophila* Computational Brain Model Reveals Sensorimotor Processing. *Nature*, 634: 210–219.

Shohat-Ophir, G.; Kaun, K. R.; Azanchi, R.; Mohammed, H.; and Heberlein, U. 2012. Sexual Deprivation Increases Ethanol Intake in *Drosophila*. *Science*, 335(6074): 1351–1355.

Shulman, C.; and Bostrom, N. 2021. Sharing the World with Digital Minds. In Clarke, S.; Zohny, H.; and Savulescu, J., eds., *Rethinking Moral Status*, 306–326. Oxford University Press.

Slashdot. 2026. Male Fruit Fly Brain Trained to Play Doom. Story and comment thread, 11 September 2026 (comments 66335190 and 66335380).

Solvi, C.; Baciadonna, L.; and Chittka, L. 2016. Unexpected Rewards Induce Dopamine-Dependent Positive Emotion-like State Changes in Bumblebees. *Science*, 353(6307): 1529–1531.

*Sorrell v. IMS Health Inc.* 2011. 564 U.S. 552.

United Kingdom. 1986. Animals (Scientific Procedures) Act 1986, c. 14, section 1.

United Kingdom. 2022. Animal Welfare (Sentience) Act 2022, c. 22, section 5.

*United States v. Richards.* 2014. 755 F.3d 269 (5th Cir. 2014).

*United States v. Stevens.* 2010. 559 U.S. 460.

*Universal City Studios, Inc. v. Corley.* 2001. 273 F.3d 429 (2d Cir. 2001).

U.S. Department of Commerce. 2026. Export Administration Regulations, 15 C.F.R. §§ 734.3(b)(3), 734.7.

Virtual Fly Brain Support Forum. 2026. [Question on Whether a Simulation Could Acquire Subjective Experience]. 15 September 2026.

Winding, M.; et al. 2023. The Connectome of an Insect Brain. *Science*, 379(6636): eadd9330.

Wissner-Gross, A. 2026. The First Multi-Behavior Brain Upload. *The Innermost Loop*, 7 March 2026. The author discloses a financial interest in Eon Systems.

Wormuth, A. 2026a. DOOMFLY. GitHub repository nftechie/doomfly.

Wormuth, A. 2026b. Stonkfly. GitHub repository nftechie/stonkfly.

Zanichelli, N.; Schons, M.; Freeman, I.; Shiu, P.; and Arkhipov, A. 2025. State of Brain Emulation Report 2025. Technical Report 2510.15745, arXiv.

Zeleznikow-Johnston, A. 2026. No, We Haven't Uploaded a Fly Yet. LessWrong, 19 March 2026.

Zer-Krispil, S.; et al. 2018. Ejaculation Induced by the Activation of Crz Neurons Is Rewarding to *Drosophila* Males. *Current Biology*, 28(9): 1445–1452.e3.

Zhou, Y.; et al. 2026. Bumblebees' Orofacial Reactions to Tastes Provide Evidence for Affective Evaluation. *Proceedings of the National Academy of Sciences*, 123(28): e2529114123.