

What are memories made of? A survey of neuroscientists on the structural basis of long-term memory

Ariel Zeleznikow-Johnston^{1†},

Emil F. Kendziorra²,

Andrew T. McKenzie^{3†}

¹ School of Psychological Sciences, Monash University, Melbourne, Australia

² European Biostasis Foundation, Riehen, Canton of Basel-Stadt, Switzerland

³ Apex Neuroscience, Salem, Oregon, USA

† Corresponding authors:

Ariel Zeleznikow-Johnston

Email: arieljz.phd@gmail.com

Andrew T. McKenzie

Email: amckenzie@apexneuro.org

Abstract:

Despite the last decade's development of optogenetic methods for artificially manipulating engrams, and subsequent claims that there is a consensus that memories are stored in ensembles of synaptic connections, it remains unclear to what degree there truly is unanimity within the neuroscientific community about the neurophysiological basis of long-term memory. We surveyed 312 neuroscientists, comprising one cohort of experts on engram research and another of general neuroscientists, to assess this community's views on how memories are stored. While 70.5% of participants agreed that long-term memories are primarily maintained by neuronal connectivity patterns and synaptic strengths, there was no clear consensus on which specific neurophysiological features or scales are critical for memory storage. Despite this, the median probability estimate that any long-term memories could potentially be extracted from a static snapshot of brain structure was around 40%, which was also the estimate for whether a successful whole brain emulation could theoretically be created from the structure of a preserved brain. When predicting the future feasibility of whole brain emulation, the median participant estimated this would be achieved for *C. elegans* around 2045, mice around 2065, and humans around 2125. Notably, neither research background nor expertise level significantly influenced views on whether memories could be extracted from brain structure alone. Our findings suggest that while most neuroscientists believe memories are stored in structural features of the brain, fundamental questions about the precise physical basis of memory storage remain unresolved. These findings have important implications for both theoretical neuroscience and the development of technologies aimed at preserving or extracting memory-related information.

Keywords

memory; engrams; expert survey; brain preservation; whole brain emulation

Introduction

Through what physical means do brains retain information for use weeks, months, or years later? Recall of long-term memories – be it recoiling from a distinct taste after a food-induced illness six months prior, or singing the lyrics of a song one hasn't heard in years – enable an animal's behaviors to be shaped by a lifetime's worth of accumulated experience. For this to be possible, earlier experiences must create some kind of 'memory trace' in an animal's brain, meaning there must be some enduring physical record that bridges the gap between an initial experience and its subsequent recall (1). But among the many potential mechanisms for information storage in the brain, which are actually responsible for long-term memory retention?

There is evidence that robust, long-lived memories are likely retained by relatively stable and static aspects of neurophysiology. Long-term memories can be recalled even after periods of prolonged global neuronal depolarization and inactivity, as can occur in deep hypothermic circulatory arrest (2,3), or despite temporary protein synthesis inhibition (4). In contrast to working and short-term memory then, it is thus unlikely that sustained, dynamic neural activation or uninterrupted protein synthesis is required for long-term memory. Instead, over the past century, neuroscientists have suggested a variety of structural candidates for the physical basis of long-term memory. A non-exhaustive list includes: synaptic strength alterations mediated by receptor insertion, removal, or subunit modification (5); synaptogenesis (6); intracellular phosphorylation, epigenetic, or other molecular modifications (7); changes in axonal myelination (8); modifications of perineuronal nets and the extracellular matrix (9); and many others. These candidates are not necessarily mutually exclusive, as long-term memory may well depend on a variety of these mechanisms and their interactions. Additionally, these mechanisms may be redundant, with multiple parallel processes capable of maintaining the same memory information even if some mechanisms become temporarily disrupted or degraded.

Breakthroughs in the past decade have seen the search for the neurophysiological substrates responsible for long-term memories, termed 'engrams', pursued with increased vigor. These advances were kicked off by the seminal 2012 publication of a method for artificially forcing the recall of a specific memory through tagging and optogenetically stimulating a subset of hippocampal neurons (10). Subsequent experiments have since demonstrated selective memory erasure (11), artificial memory formation (12), and artificial memory linkage (13), among other impressive feats. Indeed, given that ensembles of hippocampal and cortical neurons have arguably shown to be both necessary (erasure experiments) and sufficient (forced recall and insertion experiments) for memory-related behaviors, it could be argued that the physical basis of long-term memories has now been established. Indeed, in the words of many of the scientists who have furthered engram research in the past decade, "*There is a clear consensus on where the memory engram is stored—specific assemblies of synapses activated or formed during memory acquisition*" (14).

But is it really true that there is now a consensus that long-term memories are stored in synaptic ensembles? For if so, this would imply there is a common understanding among neuroscientists as to which aspects of neurophysiology are critical for long-term memory storage, and which are irrelevant. Some properties — perhaps the specific polymer subunit count of synaptic microtubules, or the exact positions of cholesterol molecules in the neuronal membrane —

would presumably be deemed below the physical scale relevant for encoding memories. At the other extreme, some macroscopic neurophysiological properties, like total hippocampal volume, would likely be considered mere epiphenomena of the mechanisms actually responsible for memory storage. But if a consensus truly exists, then neuroscientists must have some agreement that somewhere between these lower and upper bounds there lies a 'critical scale', wherein there exists the properties or mechanisms responsible for long-term memory.

To determine whether or not neuroscientists agree on where this critical scale lies, we performed a survey of both neuroscientists generally and neurophysiology-of-memory experts specifically. Expert surveys in other disciplines, such as philosophy (15,16) and artificial intelligence (17), have proven instrumental in ascertaining what is and is not consensus within an academic community. While there exists a survey of the general US population on their beliefs about how memory works (18), and there have been surveys of memory scientists on the psychological properties of memories (19), we are unaware of any previous surveys of the opinions of experts as to the mechanics of long-term memory storage.

This survey was divided into three major components. The first queried participants about whether they believed long-term memories were stored in structural aspects of neurophysiology, and if so, through which physical mechanisms and at which critical scale(s). The second explored what the implications would be of memories being stored structurally, by asking them for their probability estimates as to whether information contained in long-term memories could theoretically be extracted from static brains with ideal morphological preservation. The final section examined a theoretical test for whether long-term memories are stored structurally, by probing participants on whether readout and reinstantiation of these memory-storing neurophysiological structures in another physical form, such as through whole brain emulation (20), could someday provide proof that long-term memories are indeed stored structurally and at a particular critical scale(s).

Methods

The survey was distributed from August-October 2024 to two separate cohorts of neuroscientists: (1) those who had published research papers directly related to the neurophysiology of memory (Engram Experts); and (2) any attendees with an abstract listed in the Computational and Systems Neuroscience (COSYNE) conference booklets from 2022-2024 (General Neuroscientists). The questions were focused on their beliefs about the physical basis of memory, as well as the implications of these beliefs in various theoretically plausible scenarios. The survey and its implementation were reviewed by the Pearl Institutional Review Board and received an exemption determination (#2024-0303).

Survey questions

The survey consisted of 28 questions divided into six sections: 'Demographics', 'Structural basis of long-term memories', 'Theoretical implications of memory storage', 'Brain preservation', 'Whole brain emulation feasibility', and 'Familiarity & comfort with the topics discussed'. Most questions were mandatory for completion, except those that asked participants to optionally provide additional commentary on their responses.

Each of the main sections (i.e. all excluding 'Demographics' and 'Familiarity') were preceded by a page of information providing contextual information and definitions required for the questions that followed. The full text of the survey can be found here: <https://osf.io/agkrn>

Recruitment

We recruited the Engram Experts and General Neuroscientists cohorts separately.

Email addresses for the Engram Experts cohort were obtained in three ways: corresponding authors whose email addresses were listed in a Pubmed search for 'engram' conducted on 2024-06-20; those in a Pubmed search for '(volume electron microscopy) AND (learning OR memory)' conducted on 2024-08-01; as well as all author emails listed in the book *Engrams* (21). The Pubmed searches used an earliest cutoff date of 2012, and only results for papers written in English were examined. We manually removed papers that were clearly not relevant to neurophysiology (e.g. articles on 'immunological engrams'). After removing duplicates, this yielded 305 unique names and corresponding email addresses.

Email addresses for the General Neuroscientists cohort were extracted from the COSYNE 2024, 2023, and 2022 conference booklets. Specifically, every attendee whose email address was listed in the abstracts sections of the booklets was included. Potential participants who were already included in the Engrams Experts cohort were removed from this group. After additionally removing duplicates, this yielded 4125 unique names and corresponding email addresses.

Fielding

Participants were invited to participate in an online survey conducted via the Survey Monkey platform, delivered via private link in an email. The invitation email informed participants the survey would take approximately 20 minutes, and offered a reward on completion of \$75

(Engram Experts) or \$20 (General Neuroscientists). Payments were provided via electronic gift-cards by means of the third-party provider Tremendous.

The Engram Experts cohort was sent an initial invitation on August 29, with subsequent reminders on September 4, 9, 14, and 19. The survey remained open until September 20. Out of the 305 emails we contacted, 32 (10.5%) bounced or failed, leaving 273 functional addresses. We received 33 responses, for a response rate of 12.1%. 25 (76%) of respondents completed the entire survey, defined as providing a response to every mandatory question.

To investigate the effect of incentives on response rate, 10% of the collected emails from the General Neuroscientists cohort were initially assigned to a pilot group that were not offered a reward. The pilot invitation was on September 30. Because this unpaid pilot group had substantially lower response rates than the paid groups, we decided to roll this pilot group into the main cohort and offer a payment to all participants (including those who had initially been in the unpaid pilot group).

The General Neuroscientists cohort was sent an initial invitation on October 3, with subsequent reminders on October 8 and 14. The survey remained open until October 16. Out of the 4125 emails we contacted, 295 (7.2%) bounced or failed, leaving 3830 functional addresses. We received 279 responses, for a response rate of 7.3%. 205 (73%) of these respondents completed the entire survey.

Analysis

Data analysis was performed using R statistical software, version 4.3.1. Statistical analyses included Hartigans' dip test for multimodality, Spearman correlation analyses, Wilcoxon rank-sum tests for pairwise comparisons, and Kruskal-Wallis tests for group comparisons. Multiple comparisons in the correlation analysis were corrected using the Benjamini-Hochberg procedure with a false discovery rate of 0.05. All reported p-values are two-tailed.

Results

Beliefs on the physical basis of memory

Participants were initially asked whether they believed it would be theoretically possible, in principle, to extract the information corresponding to a non-trivial, specific, long-term memory from a static map of synaptic connectivity. This question followed after context that defined a 'static map of synaptic connectivity' to include everything down to the level of the states of individual biomolecules, and a 'non-trivial memory' by reference to examples like 'the route for navigating a maze' or 'the memory of a password'. 45.6% of participants agreed or strongly agreed that this would be possible in-principle, given sufficiently sophisticated scientific techniques, while 32.1% disagreed or strongly-disagreed (Fig. 1 a).

When given multiple, non-exclusive choices for what additional information beyond a static snapshot might be required for memory readout, respondents most commonly selected measurements of dynamically changing neuronal activity patterns (46.5%), followed by contextual information about the animal's experiences and mental states (43.0%), and information about sensory input and motor output (37.0%). Less commonly chosen were the need for measurements of chemical and electrical gradients (15.2%) and quantum-level properties (5.2%). A small percentage also selected the option that the paradigm of the question was fundamentally flawed (5.2%).

A prominent review paper on engrams published in 2016 claimed "*There is a clear consensus on where the memory engram is stored - specific assemblies of synapses activated or formed during memory acquisition*" (14). We sought to determine whether such a consensus truly existed by asking participants whether they agreed that "*The structural basis of long-term memories primarily consists of lasting changes in neuronal connectivity and ensembles of synaptic strengths, rather than in molecular or subcellular details.*" The majority (70.5%) of participants agreed or strongly-agreed (Fig. 1 b).

We also sought to assess where participants believed the 'critical scale' exists for the physical basis of memory in the brain, by asking them about whether it would be necessary to read out features ranging from the 'quantum states of biomolecules' through to 'cellular and subcellular structures (~ 500 nm resolution)' in order to extract a long-term memory from a static brain snapshot (Fig. 1 c). The proportion of responses claiming a feature needed to be measured increased monotonically as the physical scale increased, with the threshold where a majority of participants claimed 'at least some instances' needed to be measured occurring at 'Types and locations of individual biomolecules'.



Figure 1. Beliefs about the physical basis of memory. a) Responses to the question ‘To what extent do you agree with the following statement: “Given the scientific knowledge assumed in Point A existed, it would be theoretically possible to read out the information corresponding to at least one specific

non-trivial memory from a static snapshot of the structure (including biomolecules) of an organism's entire brain." b) Responses to the question *'Some neuroscientists have suggested that, while molecular and subcellular details play a role, the majority of information for long-term memories is likely physically stored in the brain at the level of neuronal connectivity patterns and ensembles of synaptic strengths (e.g. Poo et al., 2016). To what extent do you agree with the following statement: "The structural basis of long-term memories primarily consists of lasting changes in neuronal connectivity and ensembles of synaptic strengths, rather than in molecular or subcellular details."* c) Participants' views on the critical scale for memory encoding in the brain.


Implications for memory extraction from brain structure

Differing views about the physical basis of memory have different implications for the practical feasibility of potentially extracting memory-related information from brain structure. Rather than solely considering idealized 'static snapshots' of the brain, we asked participants to provide their subjective probability estimates on whether memory-related information could theoretically be extracted from brains whose structure was preserved using existing techniques, such as aldehyde-stabilized cryopreservation.

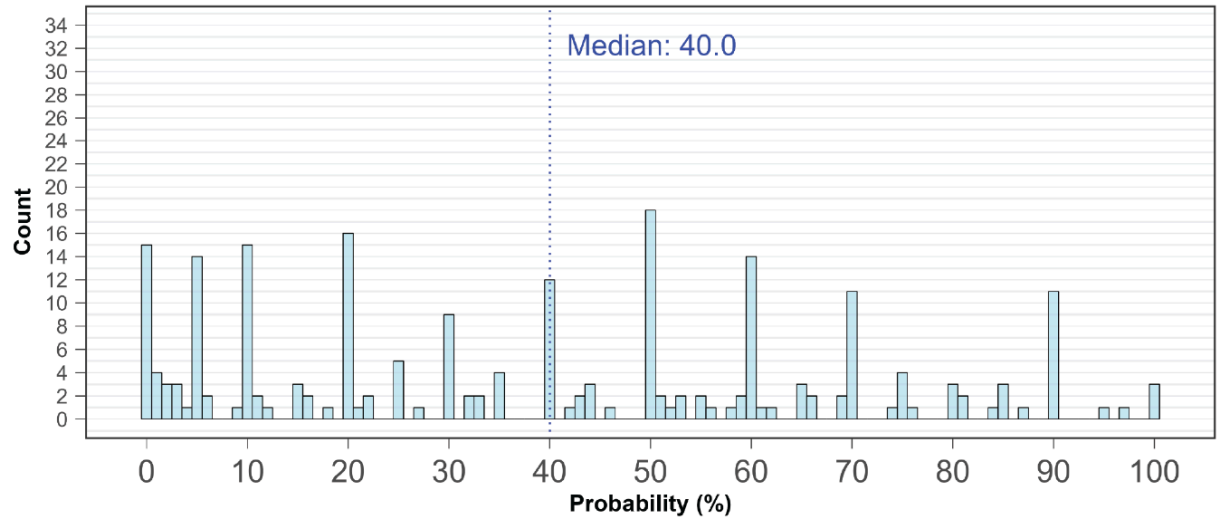
Aldehyde-stabilized cryopreservation (ASC) has been shown to preserve the brain's structural connectome, including glial cell morphology and locations (22). It is also expected to preserve the composition and spatial distribution of nucleic acids, proteins, lipids, and carbohydrates across the brain, capturing a snapshot of the brain's state just prior to the procedure (23). However, aldehyde-stabilized cryopreservation halts dynamic activity and dissipates electrochemical gradients. Were participants to believe the physical basis of memory was primarily in brain structures such as synaptic connectivity, they would presumably assign a non-trivial probability to the theoretical feasibility of extracting memory-related information from a brain preserved using this technique.

The median subjective probability estimate for the question of whether *"a brain successfully preserved with the aldehyde-stabilized cryopreservation method would retain sufficient information to theoretically decode at least some long-term memories?"* was 41.0% (Fig. 2 a). Visual inspection of the data suggested that the response distribution was bimodal, with one peak at approximately 75% and another at approximately 10%. This was supported by a Hartigans' dip test for multimodality, after a probit transformation of the data to adjust for boundary conditions of probability estimates at 0 and 100% ($D = 0.04$, $p = 0.011$).

a Subjective probability estimate that at least some memories could be decoded from a preserved brain



b Estimate that a whole brain emulation could in theory be created from a preserved brain, WITHOUT pre-recorded electrophysiological data from that animal



c Estimate that a whole brain emulation could in theory be created from a preserved brain, WITH pre-recorded electrophysiological data from that animal

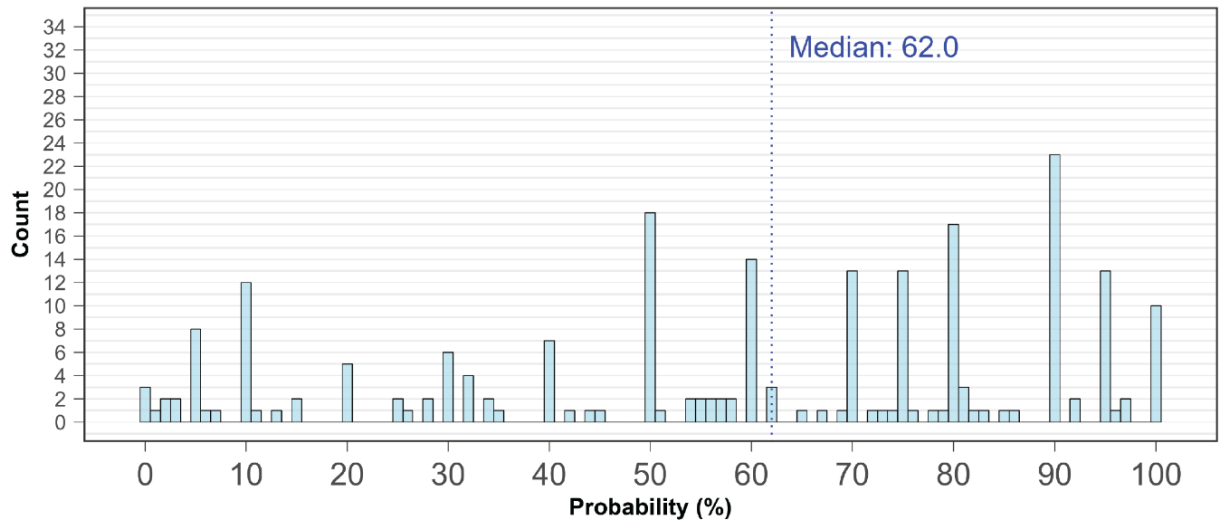


Figure 2. Subjective probability estimates for the feasibility of memory extraction from a preserved brain. a) Responses to the question ‘Given the scenario outlined above, and assuming that the preserved brain can be maintained indefinitely, what is your subjective probability estimate that a brain successfully preserved with the aldehyde-stabilized cryopreservation method would retain sufficient information to theoretically decode at least some long-term memories?’ b) Responses to the question ‘Imagine we had access to a detailed static structural map of an individual animal’s brain, such as one obtained through an advanced aldehyde-stabilized cryopreservation procedure. However, instead of having dynamic recordings from that specific individual, we only have access to a large dataset of dynamic brain activity recordings from other members of the same species (including detailed electrophysiological characterisation of all brain cell subtypes). Given this scenario, what is your subjective probability estimate that a whole brain emulation could in theory be created that accurately reproduces the long-term memory behavior of the original brain?’ c) Responses to the question ‘Assume we had access to detailed electrophysiological recordings of an individual animal’s brain cells, as well as a complete static structural map of that same animal’s brain (captured at or below the critical scale). What is your subjective probability estimate that a whole brain emulation could in theory be created that accurately reproduces the long-term memory behavior of the original brain?’

It is possible that participants may have provided a high probability to the question of whether a brain preserved “with the aldehyde-stabilized cryopreservation method would retain sufficient information to theoretically decode at least some long-term memories” even while thinking that most long-term-memory-related information would be lost in such a scenario. Accordingly, we also asked participants to provide their probability estimates on whether it would ever be theoretically possible to create a ‘whole brain emulation’ from a brain preserved using aldehyde-stabilized cryopreservation (20). We defined a whole brain emulation as “an artificial digital version of a brain ...[which] would replicate the internal causal structure of the original brain so closely that its behavioral responses to inputs and outputs would be indistinguishable from the biological version on which it was based.” As the creation of such an entity would presumably require much more long-term-memory-related information to be available than that merely required to decode at least some long-term memories, assignment of a high probability to this question would indicate strong support for the notion that long-term memories are stored in structural properties of neurophysiology.

The median probability assigned to the question of whether an emulation could theoretically be created from a preserved brain - assuming only generic knowledge of the electrophysiological properties of the neuronal subtypes for the species in question - was 40.0% (Fig. 2 b). In a different scenario, where active recordings could be taken from the brain in question prior to its preservation, the median probability increased to 62% (Fig. 2 c). Probabilities assigned in these two scenarios, as well as to the weaker question of memory decoding, were all strongly correlated (‘decoding’ to ‘emulation without pre-recording’: $\rho = 0.52$, $p < 10^{-16}$; ‘decoding’ to ‘emulation with pre-recording’: $\rho = 0.40$, $p < 10^{-10}$).

Timeline estimates for the feasibility of whole brain emulation

Even if it is theoretically possible to extract long-term-memory related information by creating an emulation of a preserved brain, there remains the question of how far from contemporary neuroscientific capabilities such a feat would be. To determine what the neuroscientific community's beliefs were on this matter, we asked participants “*by which year there will have existed a whole brain emulation of at least one member of each of the following species (C. elegans, mouse, human)*” and asked them to provide the years for which they were 10%, 50%, and 90% confident this would have been achieved. Participants were asked to round their answer to the closest available year out of ‘2025’, ‘2035’, ‘2045’, ‘2055’, ‘2065’, ‘2075’, ‘2085’, ‘2095’, ‘2125’, ‘2150’, ‘2200’, and ‘Never’.

The median participant estimate predicted a *C. elegans* would likely be emulated around 2045 [10%: 2035; 50%: 2045, 90%: 2055], a mouse around 2065 [10%: 2045; 50%: 2065, 90%: 2095], and a human around 2125 [10%: 2075; 50%: 2125, 90%: 2200] (Fig. 3).

**Estimated year by which there is an X% probability an emulation
for a member of this species will have been successfully created**

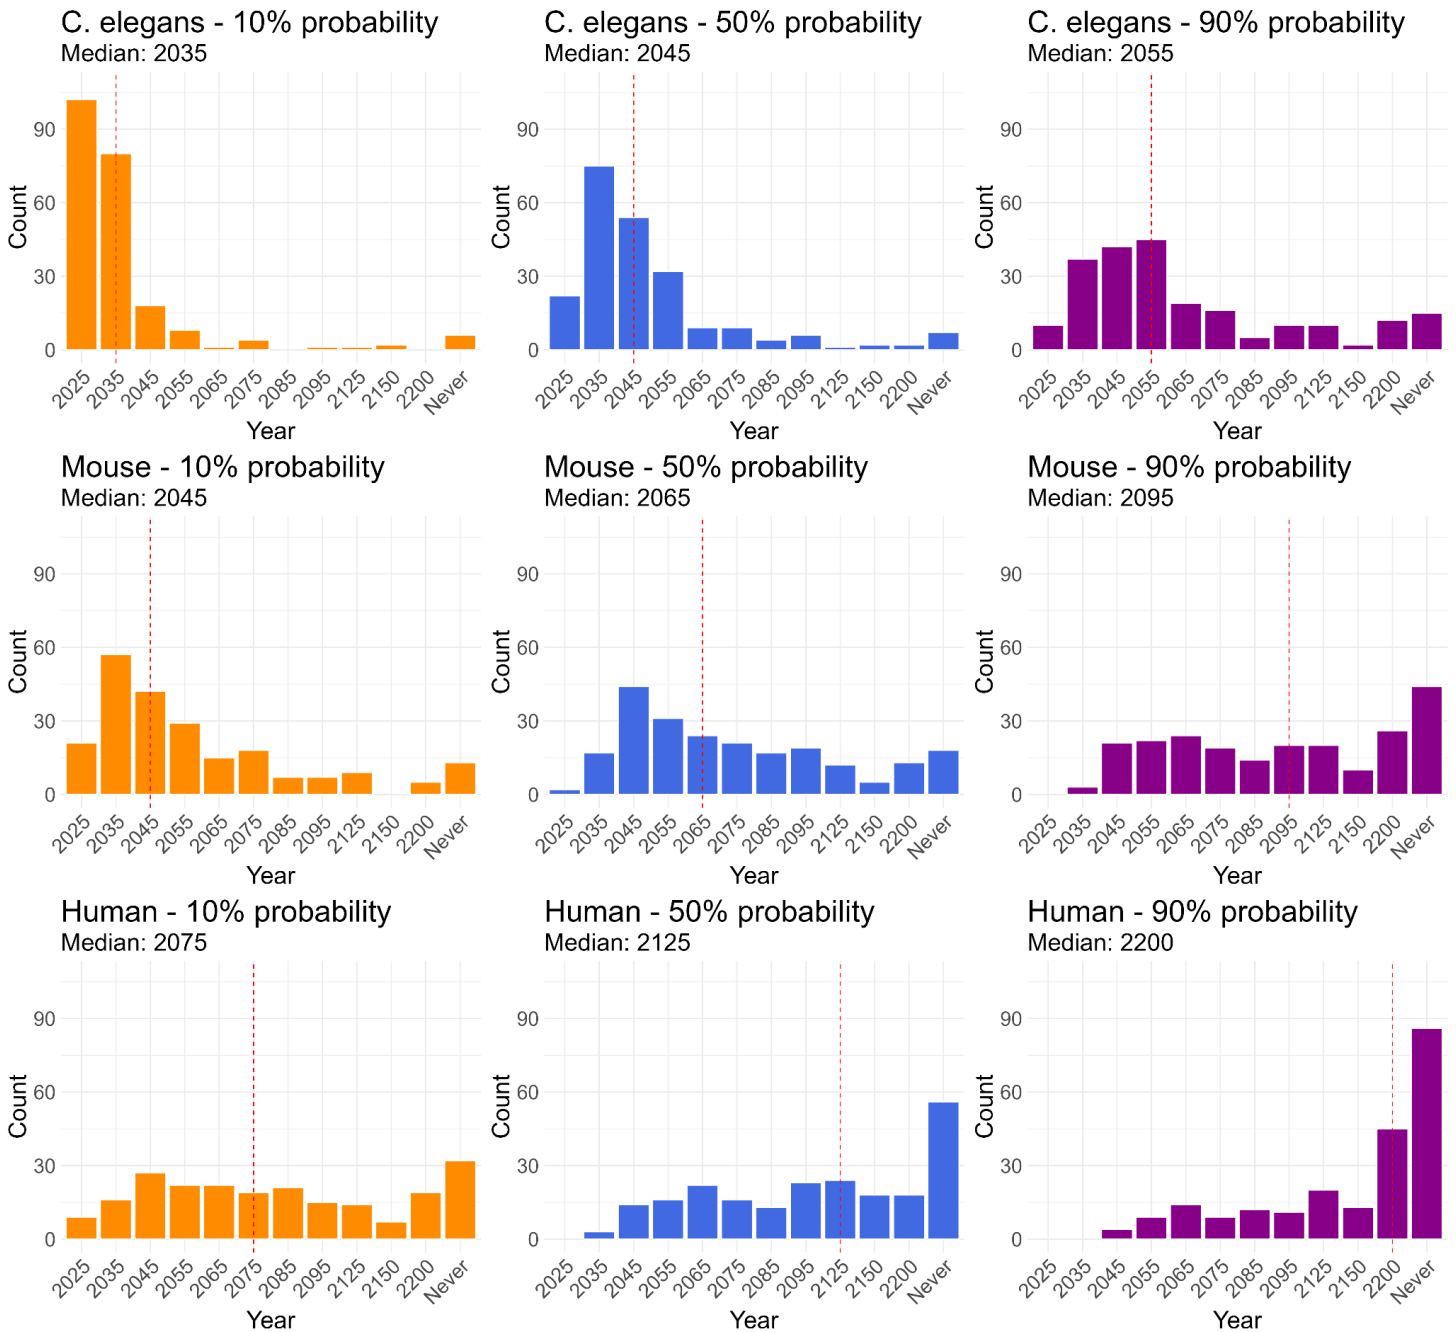


Figure 3. Estimated years by which there is an X% probability that a whole brain emulation will have been created for at least one member of the *C. elegans*, mouse, and human species.


Influence of background and expertise

We wanted to know whether experts in researching the neurophysiology of memory had different views on the physical basis of memory from those of the general neuroscientist community. Focusing on the question of whether “*a brain successfully preserved with the aldehyde-stabilized cryopreservation method would retain sufficient information to theoretically decode at least some long-term memories?*”, we found that Engram Experts did have a lower median estimate than that of General Neuroscientists (30.0% vs 50.0%), although this result was not statistically significant ($W = 2007.5$, $p = 0.10$) (Fig. 4 a).

Rather than splitting the participant data based on cohort allocation, we also sought to ascertain whether participant views differed based on their self-reported primary approach to neuroscience research (i.e. wet-lab, theoretical, both). There were no clear observed effects of research background ($\chi^2(2) = 1.55$, $p = 0.46$) (Fig. 4 b).

However, analysis of correlations between survey responses revealed several significant relationships between respondents' theoretical views and their practical predictions. Participants' estimates for the probability of being able to extract the information required for at least some long-term memories from a brain preserved using aldehyde-stabilized cryopreservation were found to be associated with their theoretical viewpoints (Fig. 5). These ASC probability estimates were strongly correlated with respondents' belief in the possibility of whole brain emulation without dynamic recordings ($\rho = 0.52$, $p < 10^{-15}$) and their belief in the theoretical possibility of memory extraction from static brain structure ($\rho = 0.47$, $p < 10^{-12}$). Demographic factors showed modest relationships, with age negatively correlated with ASC probability estimates ($\rho = -0.23$, $p < 10^{-3}$). Notably, neither memory expertise ($\rho = -0.08$, $p = 0.34$) nor preservation expertise ($\rho = -0.01$, $p = 0.96$) nor neural modeling expertise ($\rho = 0.14$, $p = 0.09$) showed significant correlations with ASC probability estimates. The strongest relationship was between participants' ASC preservation success estimates and their predictions of other participants' estimates ($\rho = 0.76$, $p < 10^{-42}$), suggesting that participants do not necessarily appreciate the wide view of beliefs that others in the community hold on this topic.

a



b

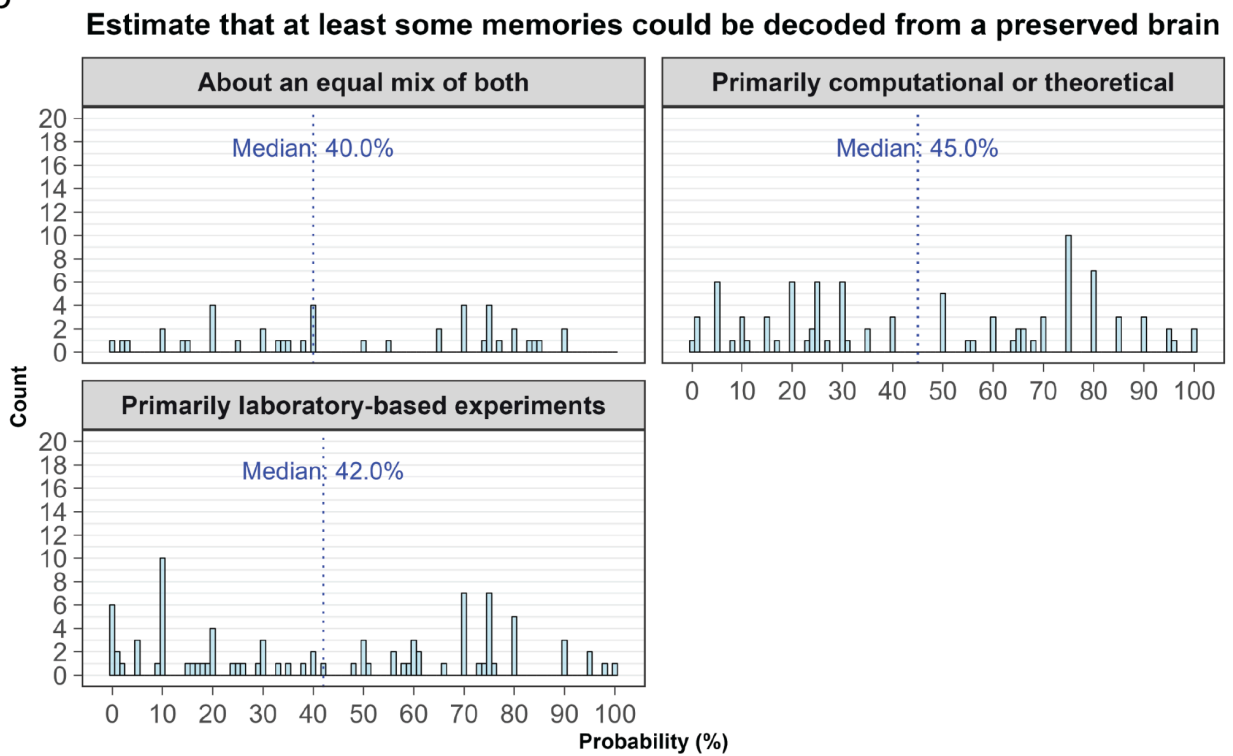


Figure 4. Subgroup analysis of subjective probability estimates for the feasibility of memory extraction from a preserved brain. a) Participant data split into either the 'Engram Expert' or 'General Neuroscientist' cohorts. b) Participant data split based on self-reported response to 'Which of the following best describes your primary approach to neuroscience research?'.

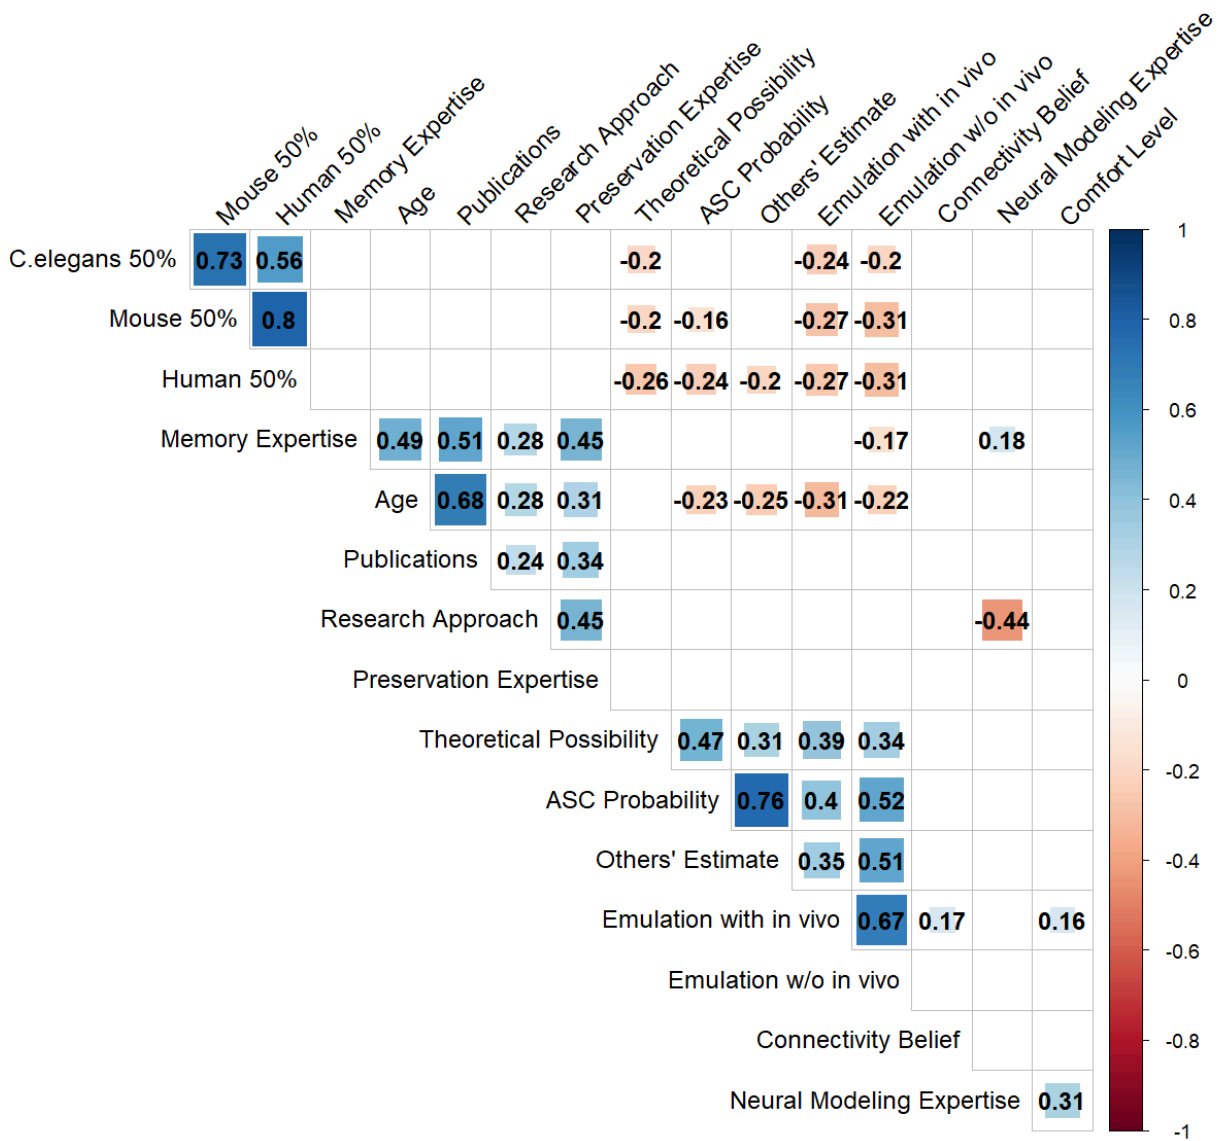


Figure 5. Correlations between participant responses to different questions. Only correlations that are statistically significant after Benjamni-Hochberg correction (false discovery rate set to 0.05) are shown. *C.elegans/Mouse/Human 50%* corresponds to the participant’s estimated year by which there is a 50% chance an emulation of that animal will have been created. *Memory expertise* corresponds to the participant’s self-reported expertise in ‘neuroscience of memory storage’. *Age* corresponds to the participant’s age. *Publications* corresponds to their number of publications. *Research approach* corresponds to their primary approach to research (wet-lab, dry-lab, both). *Preservation expertise* corresponds to their self-reported expertise in ‘brain preservation techniques’. *Theoretical possibility* corresponds to their agreement with whether long-term-memory-related information could be extracted from a static snapshot of the brain. *ASC Probability* corresponds to their subjective probability estimate that long-term memories could be decoded from a brain preserved using ASC. *Other’s estimate* corresponds to their estimate of what other neuroscientists will estimate for ASC Probability. *Emulation*

with in vivo corresponds to their probability estimate for whether a whole brain emulation could be created from a preserved brain, assuming electrophysiological recordings were taken from the animal prior to preservation. *Emulation w/o in vivo* corresponds to their probability estimate for whether a whole brain emulation could be created from a preserved brain without prior electrophysiological recordings from that animal. *Connectivity Belief* corresponds to their agreement with the statement that 'the structural basis of long-term memories primarily consists of lasting changes in neuronal connectivity...'. *Neural Modeling Expertise* corresponds to their self-reported expertise in 'Neural modeling, simulation, or emulation'. *Comfort level* corresponds to how comfortable they would feel publicly discussing these issues with colleagues.

Discussion

This study provides the first formal survey of the neuroscience community's beliefs on the physical basis of long-term memory. We found that the majority of participants endorse the claim that "*the structural basis of long-term memories primarily consists of lasting changes in neuronal connectivity and ensembles of synaptic strengths, rather than in molecular or subcellular details*". However, there was no clear consensus on exactly which neurophysiological feature or scale is critical for memory storage, with only features at or below the level of 'atomic-level features of biomolecules' being declared 'below the critical scale' by a majority of participants. Correspondingly, there were a broad range of estimates for the probability of whether long-term memories could be decoded from brain structure alone. While the median participant estimate of whether memories could be extracted from a preserved brain was 40%, there was considerable variance and some evidence for a bimodal distribution of results. Similar results were seen for the question of whether a whole-brain emulation could theoretically be created from a preserved brain, with participants estimating even odds it will be possible to create these for a *C. elegans* by 2045, a mouse by 2065, and a human by 2125. No notable effects of participant background or expertise were observed, except for a small negative correlation between self-reported expertise in 'neuroscience of memory storage' and estimates of whether a whole-brain emulation could possibly be created from a preserved brain without prior *in vivo* recordings.

We did not find a consensus on which particular neurophysiological features are critical for memory readout. Participants did indicate some boundaries, with a majority stating that the atomic-level composition of biomolecules can be ignored, while subcellular structures at ~500 nm resolution would definitely need to be ascertained. Between these though, no clear finding emerged for which particular features between these two scales are crucial for memory storage.

This is perhaps unsurprising, as despite the progress made by neuroscientists in manipulating memories over the past decade, researchers have noted that the theoretical framework for precisely identifying and understanding engrams is still absent (24). Demonstrations that long-term memories can survive periods of electrocerebral silence (2), or can be manipulated through optogenetic activation of particular neurons and erasure of particular synapses (10,11,25), provides strong evidence that memory is dependent on structural features of neurophysiology. Nonetheless, when a mouse recalls the route through a maze, it is still unclear if this representation ultimately depends on the specific densities of receptors on particular dendritic spines in particular circuits, the coarse-grained connectivity strength of several populations of neurons, or any other element or combination of neurophysiological features. Consider how in genetics, it is understood how DNA triplets (codons) encode protein sequences through the processes of transcription and translation. While memory neuroscientists agree that synapses are crucial, there is no agreement on which features (if any) are equivalent to a codon.

Despite an inability to identify which neurophysiological features are critical for long-term memory storage, a substantial proportion of participants indicated that memories could probably be decoded from the static structure of a preserved brain. This held even to the extent that the median estimate of whether 'a brain preserved by aldehyde-stabilized cryopreservation could

potentially be used to create a whole-brain emulation capable of reproducing the long-term memory behavior of an animal' was 40%. This optimistic view on what information is retrievable from structure alone may have been influenced by the recent demonstration that a whole-brain model of *Drosophila*, reconstructed from preserved brain slices imaged using electron microscopy, partially reproduced the animal's natural feeding behavior (26). While this was merely an observation of instinctual behavior, rather than individual-specific long-term memory, it is possibly suggestive of why only 10.3% of survey respondents indicated there was a less than 10% probability of a human whole-brain emulation ever being achievable. Correspondingly, as the feasibility of creating more sophisticated models of animal brains increases to the point of being able to potentially create a faithful whole brain emulation, it should become easier to test whether an emulation can indeed recreate an animal's long-term memories. If this eventually proves possible, manipulations of the emulation should further clarify which neurophysiological features are actually critical for memory.

We note that should it become possible to decode memories from a preserved brain, let alone create a whole brain emulation, this would have profound moral, philosophical, and societal consequences. As this survey focused solely on establishing the potential feasibility of this possibility, we asked participants to disregard these considerations when providing their responses. However, the substantial number of neuroscientists who indicated that memories could likely be extracted from preserved brains suggests that examining the ethical implications of this technology should not be dismissed. For example, these include the potential for life extension through means of whole brain emulation (27), or privacy violations that could occur through inappropriate memory extraction from a person's preserved brain (28). Given the median participant response indicated a 10% probability of the first human brain emulation occurring around 2075, and 50% by 2125, these prospects and pitfalls need to be understood in the decades before these capabilities might become available.

There are several limitations with both the data collection and analysis of this study. Firstly, as we received responses from only a subsample of our target populations (Engram Experts: 12.1%, General Neuroscientists: 7.3%), our findings may be subject to response bias. However, this possibility is mitigated by the fact that our invitation email only revealed this was a survey on the neuroscience of memory (i.e. it did not provide any detail of the specific questions), as well as the high completion rate (~75%) of those respondents who began the survey. Secondly, while we restricted the survey to questions about long term memory, we did not ask participants to provide separate responses for the various subtypes of long term memory. This precludes the possibility of determining whether participants have different views about the physical nature of engrams underlying episodic memory from those that form the substrate of procedural memories. Lastly, we note that our analysis of the data is incomplete, as participants provided additional information through hundreds of free-text comments that this report has not examined.

In conclusion, this survey reveals that while most neuroscientists believe long-term memories are stored in structural features of the brain, there is no clear consensus on which specific neurophysiological features or scales are critical for memory storage. The finding that a substantial proportion of participants believe memories could theoretically be extracted from preserved brains, despite the aforementioned uncertainty about the precise physical basis of

memory, suggests an underlying confidence in the structural nature of memory storage. However, the lack of agreement on critical features and high variance in probability estimates for the feasibility of memory extraction or brain emulation, indicates that fundamental questions about memory's physical basis remain unresolved. Despite this, as technological capabilities advance in the realm of whole brain emulation, we may gain empirical evidence to resolve these theoretical uncertainties. Future studies would benefit from examining different memory subtypes separately, as well as analyzing the rich qualitative data provided by neuroscience experts. Because understanding the physical basis of memory storage has consequences that extend far beyond the laboratory, the ethical and philosophical implications of our observations of the neuroscience community's beliefs should be broadly examined before they become increasingly practically relevant.

Abbreviations

ASC, Aldehyde-stabilized cryopreservation; **COSYNE**, Computational and Systems Neuroscience; **LTM**, Long term memory.

Acknowledgements

Research funding for this study was supported by a CryoDAO grant. We would like to thank Kenneth Hayworth and John Smart for helpful personal communications.

Data availability

A list of the survey questions is available here: <https://osf.io/agkrn>. The full set of participant response data is available here: <https://osf.io/bas2u>.

Author contributions

A.Z.-J., E.F.K., and A.T.M. conceptualized the study. A.Z.J. performed the data analysis and wrote the initial version of the manuscript. A.Z.J and A.T.M. performed critical review of the manuscript and helped to shape its conceptualization and content. All authors approved the final manuscript.

Competing interests

Andrew McKenzie is an employee of Oregon Brain Preservation, a non-profit brain preservation organization, and a director of Apex Neuroscience, a non-profit research organization. Emil Kendziorra is a shareholder and CEO of Tomorrow Biostasis, a biostasis provider, President of the Board of the European Biostasis Foundation, a non-profit research foundation, a shareholder and director at Oxford Cryotechnology, Inc., a cryopreservation research organization, and a board member at CryoDAO, a Swiss research association.

References

1. Josselyn SA, Köhler S, Frankland PW. Finding the engram. *Nat Rev Neurosci*. 2015 Sep;16(9):521–34.
2. Percy A, Widman S, Rizzo JA, Tranquilli M, Elefteriades JA. Deep Hypothermic Circulatory Arrest in Patients With High Cognitive Needs: Full Preservation of Cognitive Abilities. *Ann Thorac Surg*. 2009 Jan;87(1):117–23.
3. Stecker MM, Cheung AT, Pochettino A, Kent GP, Patterson T, Weiss SJ, et al. Deep hypothermic circulatory arrest: I. Effects of cooling on electroencephalogram and evoked potentials. *Ann Thorac Surg*. 2001 Jan;71(1):14–21.
4. Davis HP, Squire, Larry R. Protein Synthesis and Memory: A Review. *Psychol Bull*. 1984;96(3):518–59.
5. Lamprecht R, LeDoux J. Structural plasticity and memory. *Nat Rev Neurosci*. 2004 Jan;5(1):45–54.
6. Yang G, Pan F, Gan WB. Stably maintained dendritic spines are associated with lifelong memories. *Nature*. 2009 Dec;462(7275):920–4.
7. Giese KP, Mizuno K. The roles of protein kinases in learning and memory. *Learn Mem*. 2013 Oct 1;20(10):540–52.
8. Xin W, Chan JR. Myelin plasticity: sculpting circuits in learning and memory. *Nat Rev Neurosci*. 2020;21(12):682–94.
9. Tsien RY. Very long-term memories may be stored in the pattern of holes in the perineuronal net. *Proc Natl Acad Sci*. 2013 Jul 23;110(30):12456–61.
10. Liu X, Ramirez S, Pang PT, Puryear CB, Govindarajan A, Deisseroth K, et al. Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature*. 2012 Apr;484(7394):381–5.
11. Hayashi-Takagi A, Yagishita S, Nakamura M, Shirai F, Wu YI, Loshbaugh AL, et al. Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature*. 2015 Sep;525(7569):333–8.
12. Vetere G, Tran LM, Moberg S, Steadman PE, Restivo L, Morrison FG, et al. Memory formation in the absence of experience. *Nat Neurosci*. 2019 Jun;22(6):933–40.
13. Ohkawa N, Saitoh Y, Suzuki A, Tsujimura S, Murayama E, Kosugi S, et al. Artificial Association of Pre-stored Information to Generate a Qualitatively New Memory. *Cell Rep*. 2015 Apr 14;11(2):261–9.
14. Poo Mu-ming, Pignatelli M, Ryan TJ, Tonegawa S, Bonhoeffer T, Martin KC, et al. What is memory? The present state of the engram. *BMC Biol* [Internet]. 2016 May 19 [cited 2021 Feb 25];14. Available from: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4874022/>
15. Bourget D, Chalmers D. Philosophers on Philosophy: The 2020 PhilPapers Survey. 2021;46.
16. Francken JC, Beerendonk L, Molenaar D, Fahrenfort JJ, Kiverstein JD, Seth AK, et al. An academic survey on theoretical foundations, common assumptions and the current state of consciousness science. *Neurosci Conscious*. 2022 Aug 12;2022(1):niac011.
17. Grace K, Stewart H, Sandkühler JF, Thomas S, Weinstein-Raun B, Brauner J. Thousands of AI Authors on the Future of AI [Internet]. arXiv; 2024 [cited 2024 May 29]. Available from: <http://arxiv.org/abs/2401.02843>
18. Simons DJ, Chabris CF. What People Believe about How Memory Works: A Representative Survey of the U.S. Population. Santos L, editor. *PLoS ONE*. 2011 Aug 3;6(8):e22757.
19. Diamond NB, Armson MJ, Levine B. The Truth Is Out There: Accuracy in Recall of Verifiable Real-World Events. *Psychol Sci*. 2020;31(12).
20. Sandberg A, Bostrom N. Whole brain emulation: A roadmap. Future of Humanity Institute,

Oxford ...; 2008.

21. Gräff J, Ramirez S, editors. Engrams [Internet]. 2024. (Advances in Neurobiology). Available from: <https://doi.org/10.1007/978-3-031-62983-9>
22. McIntyre RL, Fahy GM. Aldehyde-stabilized cryopreservation. *Cryobiology*. 2015 Dec 1;71(3):448–58.
23. McKenzie AT. Glutaraldehyde: A review of its fixative effects on nucleic acids, proteins, lipids, and carbohydrates [Internet]. Open Science Framework; 2019 May [cited 2021 Jun 5]. Available from: <https://osf.io/8zd4e>
24. O’Sullivan FM, Ryan TJ. If Engrams Are the Answer, What Is the Question? In: Gräff J, Ramirez S, editors. Engrams [Internet]. Cham: Springer International Publishing; 2024 [cited 2024 Nov 4]. p. 273–302. (Advances in Neurobiology; vol. 38). Available from: https://link.springer.com/10.1007/978-3-031-62983-9_15
25. Josselyn SA, Tonegawa S. Memory engrams: Recalling the past and imagining the future. *Science*. 2020 Jan 3;367(6473):eaaw4325.
26. Shiu PK, Sterne GR, Spiller N, Franconville R, Sandoval A, Zhou J, et al. A *Drosophila* computational brain model reveals sensorimotor processing. *Nature*. 2024 Oct 3;634(8032):210–9.
27. Zeleznikow-Johnston A. *The Future Loves You: How and Why We Should Abolish Death*. Allen Lane; 2024. 368 p.
28. López-Silva P, Wajnerman-Paz A, Molnar-Gabor F. Neurotechnological Applications and the Protection of Mental Privacy: An Assessment of Risks. *Neuroethics*. 2024 May 29;17(2):31.