From Structure to Self: Philosophy of Mind as the Key to Brain Emulation

Randal A. Koene
MIMIR-IFFS Workshop, Stockholm, January 2025

Abstract

Over the past decade, data acquisition and brain mapping in neuroscience have made remarkable strides. Advanced molecular techniques now enable dynamic recording from every neuron in certain model species well-suited for calcium or voltage imaging, such as *C. Elegans*, zebrafish larvae, and *Drosophila* larvae. High-throughput electron microscopy, along with emerging methods like expansion microscopy combined with protein labeling and optical microscopy, facilitates large-scale structural data acquisition at sub-micron resolution. In some cases, such as *C. Elegans* and fruit flies, complete connectomes have been reconstructed into annotated datasets.

The primary challenge today lies in functionalizing connectome data to create working brain emulations. This "functionalization problem" depends on the intended goals of whole brain emulation (WBE), which in turn dictate the required resolution and annotations for dynamic and structural data acquisition (see: scale separation). Evaluating whether a brain emulation achieves its intended purpose requires meeting specific success criteria defined by corresponding validation metrics.

This work focuses on criteria relevant to replicating cognitive function and subjective experience in medical applications, ranging from neuroprosthetics to whole brain prostheses. Success in these areas demands demonstrating both cognitive capabilities and the preservation of a subjective sense of self. However, the criteria and validation metrics are deeply influenced by underlying assumptions about the philosophy of mind.

Philosophical presuppositions also shape the neuroethical concerns along the pathway from current technology to applied WBE. Key issues include consciousness in artificial systems, personal identity and the "copy problem," and the ethical implications of discarding imperfect emulations during model selection. Furthermore, assessing whether a WBE approach ensures survival without causing harm or suffering hinges on foundational choices in the philosophy of mind.

The concept of WBE assumes that scale separation is achievable at a suitable resolution and that psychological causal continuity can be maintained. These assumptions align broadly with functionalist perspectives, drawing on insights from Derek Parfit and Thomas Metzinger, among others.

Introduction

In the following, we aim to elucidate the importance of clear goals in the pursuit of successful whole brain emulation (WBE), to demonstrate explicit connections between those goals and assumptions based on a theory of mind, and to provide a summary layout of potential ethical issues. We begin with a brief overview of the state of the art of work towards brain emulation, and the crucial unsolved problem of system identification, particularly by means of model functionalization based on connectome data. We then present the formal *inverse problem* that underlies whole brain emulation, and how solving it relies on *scale separation*. Aiming to improve our understanding of the inverse problem involved, we present an approach to defining specific success criteria that can describe the goals of whole brain emulation in its application as a medical technology for patients with particular emphasis on a sustained subjective experience. Finally we present the tie-in to philosophical presuppositions and expected ethical considerations.

State of the art

A systematic roadmap to WBE will typically include the following sections:

1. Philosophical grounding in functionalism.
2. Specification of purpose or goals.
3. Methods for necessary brain data acquisition.
4. Methods for the reconstruction of working models to achieve successful emulation.

Additional sections may be included, for example, studies of potential embodiment following emulation, and studies of potential extensions via modification and augmentation.

The specification of goals involves producing a sufficient set of testable success criteria that must be met. For each criterion, one or more validation metrics must be found that can together determine whether that criterion has been met, and that can be feasibly applied to obtainable data and attainable models. The information processing that a brain carries out within the domain of specified goals is carried out according to certain forms of neural coding, which imposes constraints on signals and their meaningful use within the system. These constraints must result in a separation of scales within any given subsystem for processing to be substrate independent at some level, and therefore, for sensible emulation to be possible.

Acquisition of brain data is necessary at high spatial resolution where structural data is obtained, and at sufficiently high temporal resolution where dynamic data is recorded. This must be high throughput acquisition in order to deal with the vast scale of such data contained in animal and human brains. Immediately following acquisition of raw data, there are first order post-processing efforts that must be undertaken for spatial or temporal data to be useful. It is this domain, the acquisition and first-order post-processing of brain data that has progressed the most in the recent decade, and that continues to progress at an impressive rate today.

The effort to use obtained data for system identification in the reverse-engineering of a brain to successful working emulations is by comparison in its infancy. There have been a small number of recent attempts to produce useful models based on specimen-specific brain data at scale. Where the problem of system identification is addressed mainly by using morphological data, providing initial constraints by means of an identified connectome, the process is commonly called the *functionalization* problem. Any attempted functionalization must be paired with a verification and validation protocol, before implementation and application are possible with confidence in the results.

Dynamic recording

Advanced molecular techniques now enable simultaneous dynamic recording from every neuron within certain model organisms that are well-suited to calcium imaging or voltage imaging protocols, such as Caenorhabditis Elegans (Chung et al., 2013), Zebrafish (Kettunen, 2020), and Drosophila Melanogaster (Zhang et al., 2023).

In mammals, these optical methods are less practical due to the size and depth of brain tissue and accessibility issues. Nevertheless, calcium imaging has been demonstrated in mouse cortex (Zhang et al., 2023). The primary method of dynamic recording in animals makes use of electrophysiology or optoelectronic recording techniques, for example, as demonstrated for in-vivo recording with Neuropixel probes recording from hundreds or thousands of neurons simultaneously (Juavinett et al., 2019).

On the whole, high-throughput electrophysiology at high resolution and large scale is still technically very challenging and far from being a solved issue. Similarly, data recording at the molecular level remains extremely challenging at scale.

High-throughput EM and labeled ExM

By contrast, structural and morphological data collection has advanced tremendously in recent years, especially with the use of protocols for high-throughput electron microscopy (EM) and newer optical microscopy techniques that combine expansion microscopy (ExM, which relies on first expanding the sample of brain tissue between 4 and 20 times) with molecular labeling techniques, such as LICONN, a procedure combining ExM with panlabeling for protein density (Tavakoli et al., 2024).

Established high-throughput EM pipelines have produced large data sets in Drosophila (Zheng et al., 2018), using focused ion-beam scanning EM (FIB-SEM), mouse (MICrONS Consortium, 2023), using high-throughput transmission EM (TEM), and human cortex (Shapson-Coe et al., 2024). These data sets offer full volume nanoscale morphological data acquisition, supplemented with detailed segmentation, cell and synapse identification.

Emerging methods promise to improve automation - nearly eliminating the need for labor intensive human proof-reading - of segmentation and identification through sophisticated labeling techniques, such as the PRISM technology pioneered by E11 Bio.

There are now complete high-resolution data sets with associated connectome interpretations available for a few specimens of small organisms, mainly C.Elegans and the fruit fly Drosophila M.

Example of connectome reconstruction

Advances have also been made to carry out and automate the initial stages of connectomic reconstruction that are required for subsequent modeling. This post-processing of the microscopy data involves the following steps:

1. Alignment of images, which can involve reversing distortions and dealing with focus or contrast issues, noise and missing data.
2. Segmentation of images, to identify objects within the 2D images, as well as aligned 3D image stacks. This can involve resolving merge and split issues where the precise continuation of neurites from image to image can be difficult to ascertain.
3. Identification of cell types, candidate synapses and their types with corresponding ion channels, and more.
4. Human proof-reading where that is necessary to validate automated reconstruction steps and to correct errors.
5. 3D spatial reconstruction of objects and their morphology.
6. Synapse counting or synaptic vesicle counting to determine relative strengths of connections within a resulting connectome.

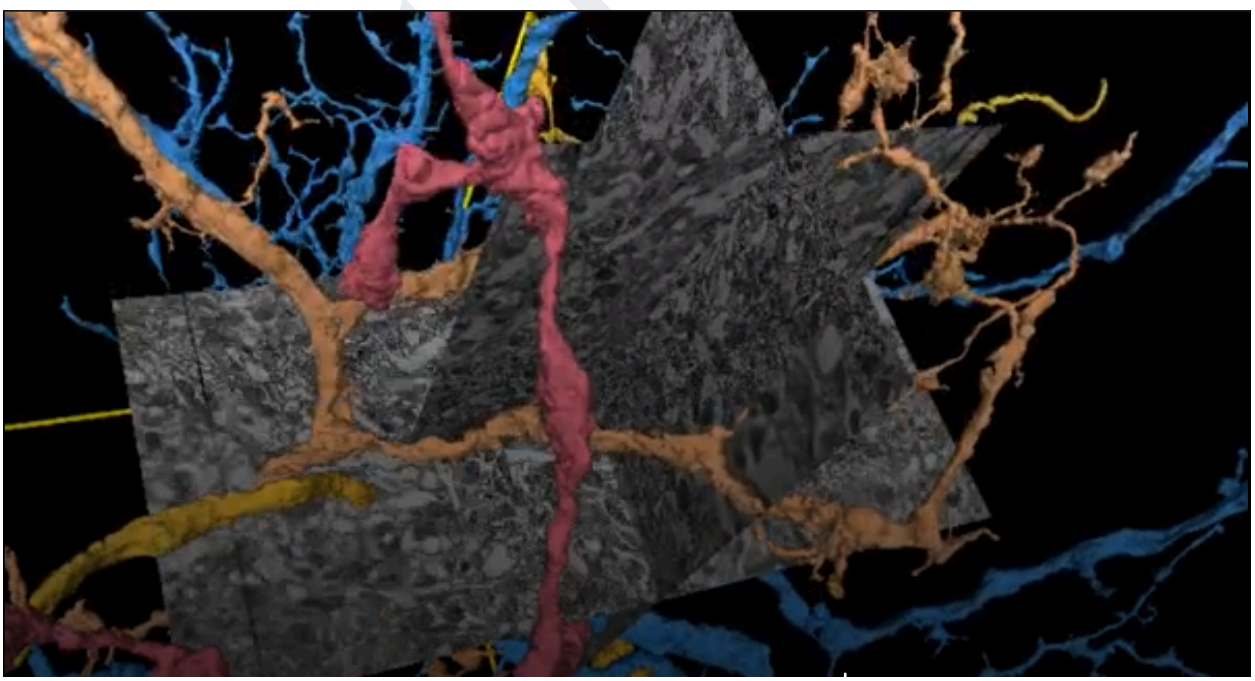

Figure 1: 3D reconstruction of selected objects from MICrONS mouse data, displayed in Neuroglancer.

The image in Fig.1 displays a typical selection of 3D reconstructed objects, such as dendrites, axons and blood vessels, as shown in the Neuroglancer tool (Jain et al., Google) for a small section of cortex of a P87 mouse in the MICrONS data set.

While impressive, the result is a static 3D reconstruction, not a runnable model or emulation of the neuronal circuitry in the brain tissue that was scanned.

Functionalizing connectomes

Given the data that has already been collected, the primary challenge today is to create a working brain emulation by successfully *functionalizing* models built using connectome and morphological brain data. As we will outline in the following sections, achieving this depends on the intended purpose or goals for which a whole brain emulation is attempted. Those goals and the consequent functionalization procedure dictate the necessary data resolution, as well as possible sample labeling protocols that can aid identification of relevant components associated with dynamic or structural data acquired (see Box 1). These requirements are imposed by a *scale separation* within the complex system, as will be explained below.

But why worry about specific goals, scale separation and data resolution associated with that at all? An alternative approach could be to set aside such concerns, which may be difficult to resolve with our current scientific understanding of information processing in the brain, and to instead endeavor to build an emulation all the way from the ground up. After all, a fundamental assumption behind the concept of whole brain emulation is the functionalist and mechanistic idea that every property and every behavior of brain and mind is an emergent result of the totality of component operations and interactions. Therefore, an emulation achieved by complete and correct model reconstructions of every component of the brain tissue, down to the molecular or atomic level, ought in principle to satisfy *every* conceivable emulation goal.

We posit that such a complete or total-depth emulation is neither practical nor desirable, and is instead merely a theoretical idealized process that can be useful in a philosophical *Gedankenexperiment* involving implications or consequences of teleportation, replication, and so forth.

Box 1: Clarification: Goals, resolutions, annotations

Goal examples:

- Emulation for the purpose of neuroscientific study. Requirements will vary depending on the study.
- Emulations used for virtualized testing of proposed neurotech devices. Requirements may include properties of brain tissue such as response to or propagation of electromagnetic stimulation.
- Emulations used during neuromedical exploration. Requirements may include inflammatory responses or reactions to chemical or biological interactions.
- Emulations for algorithmic studies that may inspire new approaches in AI.
- Emulations intended for the development of neuroprosthetic devices.

- Emulations intended to enable whole brain prosthesis or mind uploading.
- Emulations used to explore cognitive augmentation.

Examples of spatial resolutions that may be involved:
- Resolving atoms.
- Resolving molecules.
- Resolving the morphology of “compartments” used to reconstruct detailed neural cells.
- Resolving the positions of “point” neurons.
- Resolving the activity of populations of neurons.

Examples of temporal resolution or time scales that may be involved:
- The time scale of molecular interactions.
- The temporal resolution needed to discern analog neural signals.
- The time scale of phase offsets in temporal codes.
- The time scale of rhythmic brain modulations.
- Rates of neural firing. Time intervals involved approximate those of psychophysically discernable instances.

Additional annotations such as molecular labels may be needed to:
- Identify the type or determine the conductance of an ion channel.
- To determine the size or strength of a synapse, or to count synapses.
- To identify cell subtypes.
- To identify gap junctions.
- To identify or measure non-neuronal contributors, such as glial cells.

Emulation at total-depth is not feasible or desirable

Emulation involves representing operations and operands represented in one system of physical components via operations and operands in some other system of physical components. It is a conversion from one system of representations to another. Both systems involved ultimately rely on the same physical elements and laws, of course, but the *representational primitive*s of neurobiology and of electronic computation are expressed differently. Such conversion has been studied extensively in mathematical theory. In all but trivial cases, there is a *cost* associated, an *unavoidable inefficiency* that can be described and measured by calculating the Kullback-Leibler divergence, or by means of a Kantorovich-Wasserstein metric, among others (Belavkin, 2018).

Already, today’s relatively incomplete models of neurobiological circuits consume far more power to compute than their biological inspirations. It is a contrast between a human brain operating at 20-40W and compute clusters demanding KW or MW to achieve something approximating real-time performance for somewhat realistic models. It is immediately obvious that, merely on the basis of the requisite power budget, a total-depth emulation at human brain scale - let alone many human brains - is an approach that is a practical infeasibility.

Setting that aside, especially as it may be inconsequential for a first brain emulation, there is the far greater practical problem that obtaining the necessary data for a correct total-depth emulation is currently impossible. It is quite likely that it will remain infeasible or far too costly even in the long-term for the emulation of brains beyond those of very small organisms.

Even if that challenge can be overcome, there is a more fundamental problem with reliance on total-depth emulation. Doing so actively *hinders* or even prohibits the intended outcomes of brain emulation for medical and neuroprosthetic goals. To mention just a few undesired consequences:

- Reduced robustness and efficiency of the resulting emulation.
- Susceptibility to models of brain disease.
- Susceptibility to the deleterious consequences of aging.

We therefore posit that medical and neuroprosthetic goals of whole brain emulation are best served by an emulation that has the following properties:

1. *The emulation efficiently retains necessary function.*
2. *The emulation has a "footprint" that is as light as possible.*
3. *The emulation has abstracted as much of the original substrate as possible.*

Functionalization methods

Numerous methods are involved in attempts to functionalize models based primarily on morphological and connectome data. The following lists several.

It is possible to infer the likely subset to which a neural cell or a synapse belongs using knowledge of a) its location within a brain, b) its connectivity with other cells, as well as additional indicators (Lappalainen et al., 2024).

One can use established libraries describing cell and ion channel dynamic response relationships during model selection and to constrain possible parameter values (Arnaudon et al., 2023; Gouwens & Wilson, 2009; Lappalainen et al., 2024).

One can carry out direct morphological measurements, such as measuring the width of a piece of dendrite to infer capacitance and resistance values within a component, using cable model equations (Otopalik et al., 2019; Rall, 1959). Or, such as counting ready-to-release synaptic vesicles to estimate the strength of a synaptic conductance (Dürst et al., 2022; Ikeda & Bekkers, 2009; Imbrosci et al., 2022).

Knowledge of the *neural coding* strategy used in meaningful communication and processing within a specific subcircuit can provide information about functional constraints that are imposed and that can simplify modeling. For example, specific neurons may communicate through bursts of action potentials, via firing rates above some threshold rate, or may synchronize and elicit synaptic modification in a manner that is highly dependent on pre- and postsynaptic spike times (spike intervals). Groups of neurons may be connected such that they form attractors in terms of cued pattern reactivation. And in some cases, such as circuits involved in auditory sound localization, the onset time of postsynaptic activity at spatially separated synapses onto the same dendrite may be crucial to "dendritic computation" (Yamada & Kuba, 2021).

Examples of constraints that go along with coding strategies mentioned above are noise thresholds established by a signal-to-noise ratio, the necessary separation of synchronized patterns of activity by recurrent inhibition, firing synchronization modulated by rhythmic input and corresponding phase-locking of activity. Or, the activity of interacting regions may be dominated by strict phase-offsets between two communicating regions (e.g. to accomplish encoding in one region of entorhinal cortex based on retrieval from another region in entorhinal cortex). Or, guaranteed cued recall of full patterns of activity within attractors when a minimum cue of a certain size is received. Or, pattern-dependent activity within neural populations. And, of course, the model architecture itself is connectome-constrained (Lappalainen et al., 2024).

Finally, for the tuning of remaining parameters, it may be possible to use known task-specific stimulation (Lappalainen et al., 2024).

Recent Nature papers: State of the Art in functionalization

It is worth mentioning two recent Nature papers, each of which made an attempt to functionalize models based on the available electron-microscopy data set of the fruit-fly Drosophila Melanogaster.

In the paper by Shiu et al. (2024), researchers modeled and functionalized a large proportion of the Drosophila connectome using extreme simplification, yet still yielding surprisingly useful results. Every neuron in the model was represented by an identical leaky integrate-and-fire neuron. Connection weights between neurons were approximated by counting likely synapses in the connectome data. The base activity of the complete network was null, and stimulation of specific neurons was used to observe predicted downstream activations. Results from previous electrophysiology studies were able to confirm a number of the resulting predictions.

Researchers in Lappalainen et al. (2024) had a different approach, aiming instead to construct a less specimen-specific but more regular model of the Drosophila visual system for the purpose of elucidating the involvement of specific cell types within circuits forming convolutional lattices throughout successive regions of the visual system in fruit fly visual motion detection. They used connectome data to constrain the architecture of these regularly repeated, stereotypical subcircuits, as well as to estimate average connection strengths between particular cell-type pairs in those subcircuits. To establish these averaged relative weights, they counted the number of synapses for each pair of cell-types in many copies of the subcircuit that were found in the available connectome data. Using this approach, as well as information about the sign of a connection obtained from previous studies, they were able to reduce the number of free parameters in the full model to 734. These remaining parameters were optimized by visual task specific training, and the researchers were subsequently able to make cell-type specific predictions that could be confirmed experimentally.

Note that both of these attempts made use of only a small fraction of the information available in the Drosophila data set. Here is a partial list of things not done:

- No morphological measurements were made[1].
- There was no inspection of ready-to-release neurotransmitter vesicles.
- No base-level or spontaneous activity was approximated.
- There was minimal ion-channel specific differentiation.
- Synaptic spine shapes and sizes were not measured.
- No network-wide modulations were included.
- There was minimal or no subcircuit-specific variation.
- Extremely few behaviors were tested.

The inverse problem and scale separation

In system identification, the "inverse problem" refers to the process of determining the underlying system parameters or model structure based on observed input-output data, essentially trying to "reverse engineer" the system (the forward problem, control system). One infers the characteristics that produced the measured data, rather than predicting future outputs from a known system model. It is akin to trying to figure out the internal workings of a "black box" by only observing its input and output signals.

An inverse problem is said to be ill-posed if multiple system models can produce similar output data, making it difficult to uniquely identify the correct system parameters.

An inverse problem may be ill-conditioned due to noise sensitivity, meaning that the accuracy of the identified model is impacted by the corruption of real-world data by noise or variability.

As first specified by Jacque Hadamard (1865-1963), an inverse problem is well-posed for a system, $Y = f(X)$, if:

- A solution exists.
- The solution is unique.
- The solution is stable, i.e. it hardly changes when there are slight changes in initial conditions or parameters (the solution depends continuously on the data).

Nearly all well-defined scientific problems are in essence inverse problems, whether that is explicitly stated or not.

In order to even have a shot at specifying and then solving a well-posed inverse problem, one must clearly understand the problem one is seeking to solve.

[1] Aside from those involved in the decision whether an apparent candidate synaptic site should be considered an actual active synapse for the purpose of synapse counting in order to determine relative connection strengths in the connectome derived.

- What is the output or response, Y, one seeks to correctly predict, model, replicate, or emulate?
- What are the relevant inputs, stimuli or source conditions, system parameters, X, for which this needs to be done?
- What does one care about in the context of a particular set of aims?

Summarizing, to reach a successful solution, one must:

1. Define *aims* concretely.
2. Identify relevant *signals* (information, variables) that matter to those aims.
3. Properly *state the corresponding inverse problem*. Then, one can solve the inverse problem for a specific data set:
    a. Study mechanisms that operate on signals, represent, abstract and model them.
    b. Ensure that the problem is not ill-posed or ill-conditioned, is robust and generalizes adequately.

Scale separation: solving the meta problem $Y^M = f^M(X^M)$

As we will explain in sections below, the concept of scale separation is crucial both for the *forward problem (the brain)*, so that it achieves robustness, as well as for *system identification and emulation*, which is then theoretically possible and hopefully practically feasible.

The process summarized above is non-trivial and requires extensive experimentation. To reach sufficient understanding, there is a meta-problem that needs to be solved. We can call this the problem of solving $Y^M = f^M(X^M)$, to be achieved by means of some overarching scientific research protocol that is itself applied iteratively to thereby close an improvement loop that eventually results in an accurate inverse problem statement.

In this meta- control-theory or inverse-problem:

- Y^M are solved systems, described at a level of scale separation.
- X^M are usable data sets, as determined by the requirements of scale separation.
- f^M are groups of appropriate and successful analysis methods.

An effort is underway at the Carboncopies Foundation to develop a platform and an approach that makes explicit and supports a protocol of work on this meta problem, using a loop-closing system called *Generative Meta-Analysis*. A detailed introduction to this system is beyond the present scope (see Appendix).

Scale separation: reliability of complex systems

In complex systems, scale separation means that different parts of the system *operate at distinct scales* (in time, space or magnitude). This allows *separate analysis* of phenomena at each scale without needing to consider all the details of smaller scales when studying larger ones. An example in computer science would be that one can study program execution without having to consider the properties of analog transistors.

Scale separation is typically the result of a combination of system *structure* and other *constraints*. Where there is scale separation, this produces a useful protocol of signals and systems that is less prone to variability due to noise or other disturbances, one that is more reliable and robust.

It is reasonable to assume that evolutionary pressure demands that systems be *sufficiently reliable* for the duration of the survival life-span of a specific organism. An animal's survival depends greatly on reliable processing of *sensory input*, on reliable *recognition of threats and opportunities*, on reliable *learning*, and on taking reliably *beneficial response actions*. A chaotic analog system is by default prone to unpredictable random walks in terms of its states and behaviors.

It is worth noting that a small amount of randomness can incur a degree of benefit to learning in that it can be a means of escaping local optima in the search for better solutions. Nevertheless, wild random walks are unreliable and dangerous.

Scale separation: evolution, engineering, emulation

Unsurprisingly, we find in the brain that there is a large degree of systematic architecture, repeated occurrence of templated circuitry, and an application of robust signal protocols that are enforced by constraint mechanisms within each protocol.

Similarly, when engineers design robust and reliable computing while using naturally analog components that must operate in a variable environment, they also impose scale separation. For example, they do this by specifying voltage thresholds, by including parity bits, by imposing clock cycles, and much more.

Robustness requirements demand scale separation. And this must be true for brains.

Just as scale separation is vital for the original system, it is also essential for successful emulation. It allows us to replace parts of the substrate with mechanisms implemented in a different substrate. To do so, one must find the relevant scale separation at levels of the pyramid of mechanisms upon mechanisms found in nature. Above identified levels of scale separation, that which we wish to emulate is then *substrate independent*.

Identifying points of scale separation in the context of goal-directed success criteria, within various functional subsystems of a brain, can enable one to clearly express the inverse problem that is to be solved. One can consequently explore those methods of analysis that may be able to find its solution.

Success criteria for subjective experience

Before goals have been clearly defined in terms of success criteria, we may loosely describe the purpose of whole brain emulation for medical or neuroprosthetic application as producing an emulation with which one achieves satisfactory survival by (some consensus of) the standards of patients. We may commence a further exploration, noting that these goals appear to focus on cognitive processing as it is experienced.

At the Carboncopies Foundation, we made an attempt to draft a list of possible success criteria through a process involving three surveys working with interested parties. We applied a protocol with the abbreviation ICEDA, which was repeated as many times as necessary:

1. **Identify** candidate success criteria.
2. **Categorize** and **Reduce** the batch of candidates.
3. **Evaluate** candidates, rejecting or keeping them based on requirements such as whether they can be met by satisfying objective metrics.
4. **Derive** corresponding sets of metrics with which to validate if success criteria have been met.
5. **Apply** the resulting success criteria to attempts at brain emulation.

The full protocol, the collected survey data, and several resulting categorizations and reductions are available in separate publications. Here, we summarize to give a sense of our draft results.

Categories in the draft set were:

- **Psychological**: behavior & responses, personality traits, experience+feelings+beliefs.
- **Cognitive**: capabilities & performance, awareness, problem solving & reasoning, learning+encoding+retrieval of memories.
- **Neural Circuits**: components & properties, I/O within an acceptable envelope.
- **Embodiment**: sensation, environmental connection, implementation & safeguards.
- **Societal**: recognition & acceptance by others.
- **Other**: protocol, process development, legal & regulatory.

Following categorization, reduction and evaluation of candidate success criteria, the surveys resulted in a total of 25 proposed success criteria. Of those, 12 satisfied requirements that were needed to be usable during initial research, because such research typically involves pieces of brain tissue or small animal brains, rather than responsive human patients. The one-line summary descriptions and success criteria ID code of these 12 are shown in Box 2.

Box 2: Success criteria from surveys for the development of validation metrics usable in

initial research.

P-3 Observable behavior and responses are sufficiently similar, similar Illusions.

C-3 Mental performance is at least as fast.
C-4 Dynamic evolution through plasticity.
C-8 A specified set of memories, possibly encoded specifically for the test, is retrieved sufficiently well.
C-11 Objectively verifiable memories are retrieved with a sufficiently similar set of correct and incorrect answers.
C-13 Able to form new memories and learn new skills at least as well as before.

N-1 Reconstruction of neuronal circuits through system identification and tuning of properties is sufficiently accurate.
N-2 MIMO neural traces match within a specified envelope.
N-3 Spatio-temporal patterns, brain rhythms, synchronization are sufficiently similar when stimulated and in specific brain states or modes.
N-4 Circuit function breaks in similar ways & similar robustness

E-1 Sensory input for sensory experience provided.
E-3 Being able to act or interact, affecting the world.

Success criteria determine requirements for data and translation to parameters

System identification involves finding parameter values for the mathematical representation of brain processes within models (component simulations) that compose a runnable emulation. The mathematical representations must suit and be able to satisfy our collection of success criteria. There may be high-level representations that are fairly abstract, low-level descriptions closely approximating neurobiological mechanisms, and representations anywhere on the spectrum in-between, so that they respect the relevant scale separation.

There should be a set of possible representations that together satisfy a particular set of success criteria. The representations in that set will have a multitude of parameters. If well-chosen, then the parameters may be quantified or tuned using data that is obtainable from a biological sample. One must find a *solvable relationship* between the *observable (and observed) data* and the *parameters that need to be determined*.

When using structural (morphology and connectome) data, the system identification problem is addressed by a) deriving structure at an appropriate resolution, and b) solving the functionalization of corresponding models.

We can summarize the process of mapping success criteria to data requirements as follows:

1. You have a proposed description or representation.
2. There is a corresponding identified structure.

3. You have a proposed process for model functionalization.
4. There needs to be data to make both (structure and functionalization) possible.
5. You find types of *measurable data that are well-correlated* with both.
6. These are then the useful and necessary data to collect.
7. Consequently: *Exploring the right data acquisition depends critically on success criteria*.

Philosophical presuppositions

A fundamental underpinning of whole brain emulation is a form of *functionalism*. Requisite presuppositions include:

- The transformation of signals, or the processing of information, and the operations or functions involved in that are the core requisites of mental experience and of subjective experience.

- *Any* sufficient set of functions that carry out the requisite signal transformations or information processing will result in the same mental experience and the same subjective experience - assuming a correct identification of *scale separation* for the particular goals of emulation.

- Replacing individual signal carriers or individual processing mechanisms in a manner that does not alter downstream processing makes no difference to the outcome.

- There is no additional attribute or property of the physical or material constituents involved that matter to the resulting cognitive experience, assuming that one has correctly identified the relevant signal transformations or information processing that are achieved by physical mechanisms.

On the whole, this position is in alignment with arguments by Derek Parfit (1984) and other philosophers taking a strong functionalist stance.

Replicating cognitive capabilities and preserving sense-of-self

In medical applications, from neuroprosthesis to whole brain prosthesis, success depends on demonstrating both a) preserved cognitive capabilities, and b) a preservation of a subjective sense-of-self.

Replicating one's cognitive capabilities refers to the abstraction and reimplementation of a collection of experienced mental faculties that matter to us in our daily lives. These are the faculties that provide the information from which we build our internal world-model simulation (Metzinger, 2004), and which allow us to act upon that. This can involve the actual world-model simulation, a putative world-model simulation such as is needed for goal-directed planning, world-models of imagining, of dreaming, or of day-dreaming. A portion of that world-model we self-identify with, possibly because it is the portion for which we have learned that we exert direct control or that we reliably experience in all circumstances.

In reference to a, similarly, internally generated *phenomenal Now*, that is a representation of a most recently modeled interval (not coincidentally closely correlated with the temporal granularity of working memory), these cognitive processes enable us to carry out an internal validation that confirms subjective experience. At that level of the internal model, we may test the state of variables (established by information processed *a-priori*), for example by asking, “Am I aware right now?” (Blackmore, 2017). If our answer to the question is perpetually “yes”, then the sequence of affirmative sample points forms a temporal axis along which we are continuously certain that we subjectively exist, as if in a continuous stream[2] (Fig.2). This, even though the validation test carried out does not rely in any way on a separate homunculus or a further fact, but instead on the processed and stored state of a world-model and its self-model subset that are the inevitable product of the cognitive processes carried out in our brain.

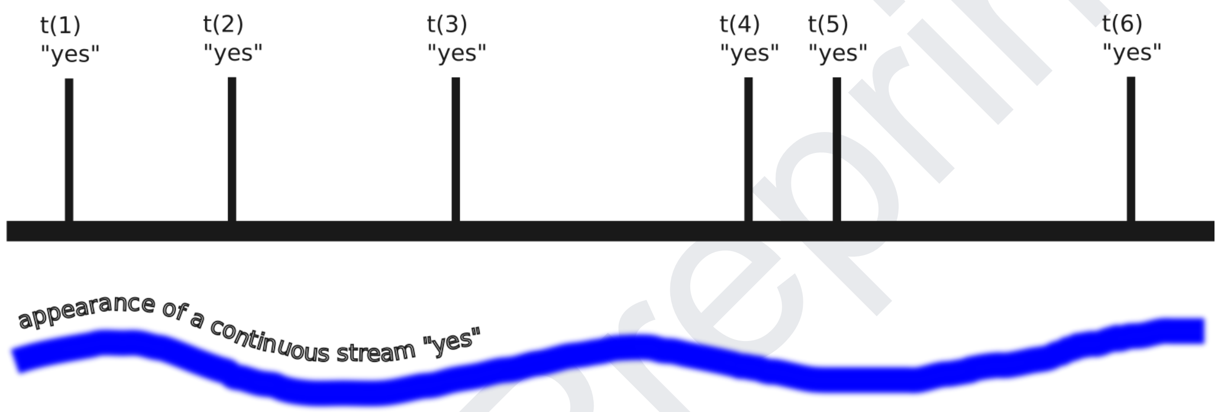

Figure 2: The appearance of a continuous stream of subjective awareness as a consequence of successive judgments based on a self-model and a phenomenal Now.

Additionally, and very importantly, we also recognize ourselves uniquely through characteristics that are defined by:

- Our unique memories. Here, memories refer to all types of representations that contain some form of dependence on the past, i.e. a hysteresis in their mathematical functional form. This may extend beyond the types of memory typically referred to in cognitive science, such as episodic memory, implicit procedural memory, explicit declarative memory, semantic memory, sensory buffers, and muscle memory.
- Our motivating drives and preferences. In a sense, many of these are themselves informed by our memories, through which we filter perception and action, though they will also depend on a personal blue-print bestowed by genetics, the developmental layout of our specific brain, and a unique balance of the modulatory chemistry (e.g. hormones).

[2] An interesting corollary is the peculiar relationship to reality caused by sample-based “experience”. It is often said that, because we are macroscopic beings, we experience the world as continuous, while quantum theory points to discrete, quantized units. But, the evidence, in the form of reflective samples, that our minds can use from which to infer a continuous experience is itself discrete, and at a far lower temporal resolution.

- Habitual interactions with other parts of our world-model, such as what we know of others, our personal social associations.

Theoretical focus: The category of Metzinger, Parfit, et al.

In picking the philosophical approach that is most appropriate and useful for the study of brain emulation, why do we focus on insights by Thomas Metzinger, Derek Parfit, Susan Blackmore, Daniel Dennett, Michael Graziano, and others with backgrounds in philosophy or neuroscience who generally adhere to a similar analytical approach? Here, we could discuss why many alternative perspectives are unlikely to survive closer scrutiny or are unhelpful, but instead we will explain the particular usefulness of our preferred class of insights.

Each of the authors mentioned takes what is essentially a *functionalist* perspective. Not only that, they refuse to abandon ship prematurely and maintain a *reductionist* approach to the study of consciousness and subjective experience, which has been an exceedingly successful approach throughout science. Their deliberate process focuses on the characteristics and prerequisites of consciousness and of subjective experience, and they address the most relevant questions as to *how* it can be that our experience is as it is experienced. They do this in a highly structured and systematic way, which retains the hope that consequent explanations can elucidate the precise *nature of experience*, *without the need for a further fact*.

Leaning into the diligent work by Metzinger (2004), the following are the most important concepts for a conscious mind that brings about what must result in self-referential judgments of subjective experience:

- The *transparency* of underlying processing, in that, at the level of conscious deliberation one does not have insights or access to the vast network of processes that were involved in constructing world- and self-models.
- That there are *specific* and *separate* types of processes that support conscious operations.
    - This insight is supported by empirical evidence, where active *working memory* is highly correlated with states of conscious experience. From a neuroscientific perspective, it is important to note that working memory does indeed appear to depend on a very specific and separate set of mechanisms (Baars & Franklin, 2003; Koene & Hasselmo, 2007a, 2007b). These mechanisms allow for sustained, repeated, temporally ordered patterns of activity.
    - And these working memory specific forms of sustained activation *ride on top of and temporally separated from the flow of synapse-dependent network activity of structural processing*. Notably, the underlying network activity does not depend on the same rhythms of highly structured temporal processing that characterize working memory.

    - These separate processes are momentary, fleeting, and absent when unconscious, precisely as one would expect of a system that is strongly involved with conscious processes.

- That there are internal simulations constituting a *world-model*, and within that, an identified *self-model*.

- That the perception of time is itself a part of the constructed simulation, and specifically, that the brain generates world-model states that are experienced as the *phenomenal Now*. In terms of neuroscientific evidence, this can again be tied to the temporal granularity of working memory (D'Esposito & Postle, 2015; Ji et al., 2024; Storm et al., 2024).

- All conscious judgments are *post-hoc* in that they rely entirely on information that is already available in "conscious space" provided by the world-model and the phenomenal Now.

- We *learn* to recognize the self-model within the phenomenal Now as our subjective selves. Aligned with this notion, consider the developmental stage of childhood where a child switches from third person statements (e.g. "Emma wants cake") to first person statements (e.g. "I want cake").

- If we ask ourselves, "am I here, am I aware now" (Blackmore, 2017), given the proposed mechanisms above, the answer will normally be "*yes*" in that moment where a phenomenal Now has been generated and an identifiable self-model is present. Such a validation becomes an evidentiary sample point. Strung together, sample points produced by references to data in our present self-model must inevitably result in the *appearance of an experienced continuous stream* (Fig.2). Combined with the transparency of the vast underlying sea of processing, a seemingly separate subjective experience is confirmed.

Ethical considerations

At the Carboncopies Foundation, we decided that a useful approach to the ethics of whole brain emulation is to consider matters that concern:

- WBE in the context of society.
- WBE in the context of patients and animal subjects.
- WBE in the context of the responsibilities of researchers.

Furthermore, we take into consideration matters of prevailing concern, a) in the context of actions taken during the process of development towards achieving WBE, and b) in the context of a future situation where practical WBE is a technical reality.

At the present stage, it is too early to present a careful exploration of each possible issue or to propose worked out solutions. Here, we provide a list of concerns, shown in Table 1.

Fundamental issues
Consciousness and consequent potential for suffering in artificial systems.
The nature of personal identity and how that impacts the so-called “copy-problem”.
Existential risk arguments for WBE: Can WBE offset or counter X-risks posed by AI?
Existential risk arguments against WBE: Does WBE itself contribute to X-risk?
Should WBE technology be developed in the spirit of open-source or closed-source?
Research phase
The treatment of animals used during study.
The treatment of (partial) WBE of mice, non-human primates. Consider that one may have to construct N separate models of any one brain before the typical “model selection” phase of any estimation procedure (as demonstrated in typical machine learning studies). What happens to the N-1 less performant models?
Dissemination of data and knowledge obtained from WBE studies. What about the effect such data and knowledge may have on the development of AI?
WBE process issues with regard to the patient
Ensuring equitable or equal access to (eligibility for) WBE technology and corresponding medical procedures. Is there a cost to the patient?
What if some nations forbid WBE and its citizens are too poor to travel elsewhere for access?
How are patients prioritized? By age, by “at risk” status, by wealth?
Well-informed patients: Do they understand risks and consequences of WBE?
Consent and privacy: The “uploading” process involves intimate access to one’s personal thoughts and memories. In fact, such access may be essential for necessary validation procedures.
The risk that providers may prey on customers by offering discounts in exchange for giving up some rights, e.g. through “terms and conditions”.
Validation of WBE and guarantees of outcomes.
During potentially necessary iterative validation tests of WBE and conscious experience, is backing out of unsuccessful results ethically possible if experienced states are essentially those of a sentient person?
What are the ethical implications of discarding imperfect emulations during model selection?

If collected data is incomplete or corrupted in some way, should one still carry out WBE in order to conform with wishes of the patient?
What if all of the “success criteria” are not met, does one keep the WBE anyway?
What to do with degraded “uploads”?
Post-upload issues with regard to persons
In the case of an elderly brain, are modifications intended to “refresh” that brain permissible, if such modifications in essence change the person?
If an “uploaded” patient becomes unhappy with their condition, should one terminate the WBE or roll back to a previous happy state?
What about the possible need for ongoing monitoring (insight and access to private mental states) as part of psychological health care for the well-being of the person?
General access to brain data for mental health-care.
Cognitive liberty: Mental self-determination, cognitive sovereignty.
Rights to types of embodiment.
The possible exploitation of “uploads”.
The legal status of “uploads”. For example, can they enter into contracts, own property, etc?
What if future insights lead to update of ethical principles after there are already emulations running? Should those new principles apply to them or does one hold to their original contracts even in light of new insights?
How about the sharing of brain data or merging of brain data?
Is the patient the only one who decides how long they may live?
If multiple copies are made illegally, i.e. against a patient’s wishes, what happens to those copies if, by then, each one has become its own person?
What if a patient chooses to make “dumbed down” copies or partial extracts of themselves to share or lease out for use with the intention to subsequently terminate them?
Is it ethical to have multiple tiers of emulation, e.g. different speeds or different augmentations available?
If there are ongoing expenses attached to an “uploaded” existence and the “upload” stops paying, is deactivation murder?
Social issues when WBE exists
What should be the roll-out procedure for WBE in order to ameliorate sudden societal or economic

impacts?
How are social dynamics impacted, such as job competition, or family relationships?
Is there a risk of entrenchment of fixed hierarchies and consequent status-quo preserving (conservative) tendencies and enhanced power concentration in a world with WBE?
How do very long lives restructure society?
Should there be permission to self-modify, to augment, to carry out personality changes? Should there be any limitations?
Should there be any constraints on the types of permissible divergences (applied changes), if those can have effects on others?
What about the ownership of WBE centers, data centers? Private, public, other?
What if a provider goes bankrupt? What happens to the "uploads" affected?
What should be the ownership status and access rights to backups?
Can emulations have children?
If there are restrictions on drug use, should there be restrictions on virtual drug use?
If copies are made that live on to become their own persons then what are the consequences to voting rights?

Table 1: A list of ethical concerns, categorized.

Issues raised for WBE and AI should feed back into improved treatment of all sentient creatures and humans

An interesting consequence of considering a world with whole brain emulation is that it leads us to consider with crisp clarity some issues that are otherwise simply taken for granted. It is reasonable to ask that in tackling these issues, resulting in improved ethical insights, some of those insights should propagate back into broader ethics. The following are just two examples.

<u>Example 1</u>: Ownership and privacy of mental processes and mental content. Compare this with the current acceptance of:

- Advertising that is intended to change one's perceptions and behavior.
- Religious indoctrination, which is presently considered a fundamental right, even though it explicitly endorses the sculpting of individual thinking, especially in children. Are the minds of children the property of their religious community?
- Learned magical thinking, spreading non-rational modes of thought with the intent to change minds and create a community or following.

- Parenting choices that instill particular values, present particular examples of behavior, quite explicitly imprinting on children. Again, this is presently considered a highly valued right of parents, even though this clearly and directly takes ownership of and modifies the minds of children. Are the minds of children property of the parents?

Example 2: Increased entrenchment of existing hierarchies and their self-preserving tendencies:

- There is already an almost unavoidable trend towards government for and by the powerful.
- A lot of power presently exists within a limited number of highly-merged conglomerate companies.
- Maintaining a status-quo is already an existing and somewhat accepted state of affairs.

Risks and realism: ASI and WBE

There are often situations where artificial (super) intelligence (ASI) is compared with whole brain emulation. Depending on the assumptions involved, this can lead to judgements such as the proposal that WBE would never have a chance to rise to the level of intelligence that is presumed to be rapidly achievable by AI. Sometimes this even leads to conclusions such as that developing WBE might be pointless, as it could never catch up with the development of ASI. While this cannot be ruled out, it is worthwhile to consider the underlying assumptions more closely.

What is intelligence actually, and more importantly, what does the effective application of intelligence depend on? After all, there are numerous anecdotes of supposedly high-IQ individuals who end up spending their lives in relatively mediocre vocations without producing any world-altering output. It is worth noting that progress is not primarily driven by introspection. Typically, significant progress in any field requires data collection and experimentation. These things are not simply computable and subject to simulation *a-priori*. Instead they depend crucially on the rate of events in our universe.

Taking this into account, we may postulate that acceleration of progress may be absolute rate limited, both for humans now, for WBE, and for ASI. Nevertheless, there are at present important differences between the characteristic constitution of human or WBE mental processes vs ASI mental processes. WBE, being based on the human brain, has a far more heterogeneous structure than the typical structure of AI algorithms. It is probably a fair assumption that a more homogenous structure relying on only a few optimizing algorithms lends itself to more straightforward and potentially more easily accelerated growth.

Then again, the heterogeneous structure of the human brain has evolved specifically to address multiple requirements of the world around us. It is equally fair to assume that there is not yet a conclusive demonstration as to whether homogeneous or heterogeneous brains are ultimately superior in all of the ways that matter.

If it turns out that there is an absolute rate limitation imposed by the collection of signals and data from the universe, through experiment or otherwise, then one can imagine protocols through which even heterogeneous brains can be brought to achieve that same rate. It is possible that the resources needed to bring heterogeneous brains to that cap may be different or greater than those needed for homogeneous brains. This is a direction of research where much future work remains to be done.

Ethical issues of not achieving whole brain emulation

Finally, we should also take seriously the ethical issues that may be involved with a decision *not* to develop and achieve whole brain emulation.

Without WBE, the existence of human persons is constrained to biologically habitable environments. Those that are livable for human mammals. In the absence of substrate independence, humans may be limited to life on a single planet and a few fragile artificially constructed habitable biosphere bubbles elsewhere in our solar system. Another way to put this is that it is almost unimaginable that a species (e.g. an alien species) that has become truly space faring, that spreads through the cosmos, would do so without achieving substrate independence and the far greater flexibility to adjust to the conditions of local environments that comes with it. There is a reason why robots are today exploring the fringes of our solar system and beyond.

Existing on a single planet is inextricably linked with a panoply of obvious existential risks. There is the ultimate fate as our sun enters its next phase. Long before that, plenty of opportunities for extinction-level cosmic impacts with the Earth. And of course, risks to our environment that emerge from our own actions.

Aside from this species-wide consideration, not achieving whole brain emulation also means that one is left with the daily risk to individual lives, such as we are accustomed to in the form of incidents, disease and aging.

Conclusion

The pursuit of whole brain emulation represents a profound convergence of neuroscience, philosophy, and technology, with the potential to revolutionize medical applications, cognitive science, and our understanding of consciousness. As we outlined, the functionalization of connectome data is the central challenge in achieving WBE, requiring not only advanced data acquisition and modeling techniques but also a clear articulation of success criteria tied to specific goals. These goals, whether focused on cognitive replication, subjective experience, or neuroprosthetic applications, are deeply influenced by philosophical presuppositions about the nature of mind and consciousness. Functionalist perspectives provide a robust framework for understanding how mental processes and subjective experience can be preserved in emulated systems.

However, the path to WBE is fraught with technical, ethical, and philosophical complexities. The functionalization problem demands innovative solutions to the inverse problem of system identification, leveraging scale separation to create efficient, robust, and substrate-independent emulations. At the same time, the ethical implications of WBE - ranging from questions of personal identity and consciousness in artificial systems to societal impacts and existential risks - require careful consideration and proactive engagement.

As we advance toward the realization of WBE, it is crucial to maintain a dialogue between scientific progress and philosophical reflection. The insights gained from this interdisciplinary effort will not only inform the development of WBE but also deepen our understanding of the human mind and its place in the universe. Moreover, the ethical frameworks developed for WBE have the potential to influence broader societal norms, particularly in areas such as cognitive liberty, privacy, and the treatment of sentient beings.

In conclusion, while the technical challenges of WBE are formidable, the potential benefits - from medical breakthroughs to the preservation of individual identity and the expansion of human existence beyond biological constraints - are immense. The journey toward WBE is as much about understanding ourselves as it is about advancing technology, and it is a journey that promises to reshape the very fabric of what it means to be human.

Appendix: Generative Meta-Analysis

Generative Meta-Analysis provides a systematic protocol that aims to accelerate the necessary improvement loop for analysis methods (Fig.3) that are applied to the discovery of valuable subsystems within large, complex nonlinear systems, by supporting ground-truth based testing and discoveries in scale separation. Testing of analysis methods is done by generating many example systems in-silico, ensuring a fully-known ground-truth that may otherwise be unobtainable. Examples are created at multiple degrees of difficulty, and variables included can be controlled precisely. The following is a brief description of the protocol for which our platform is designed.

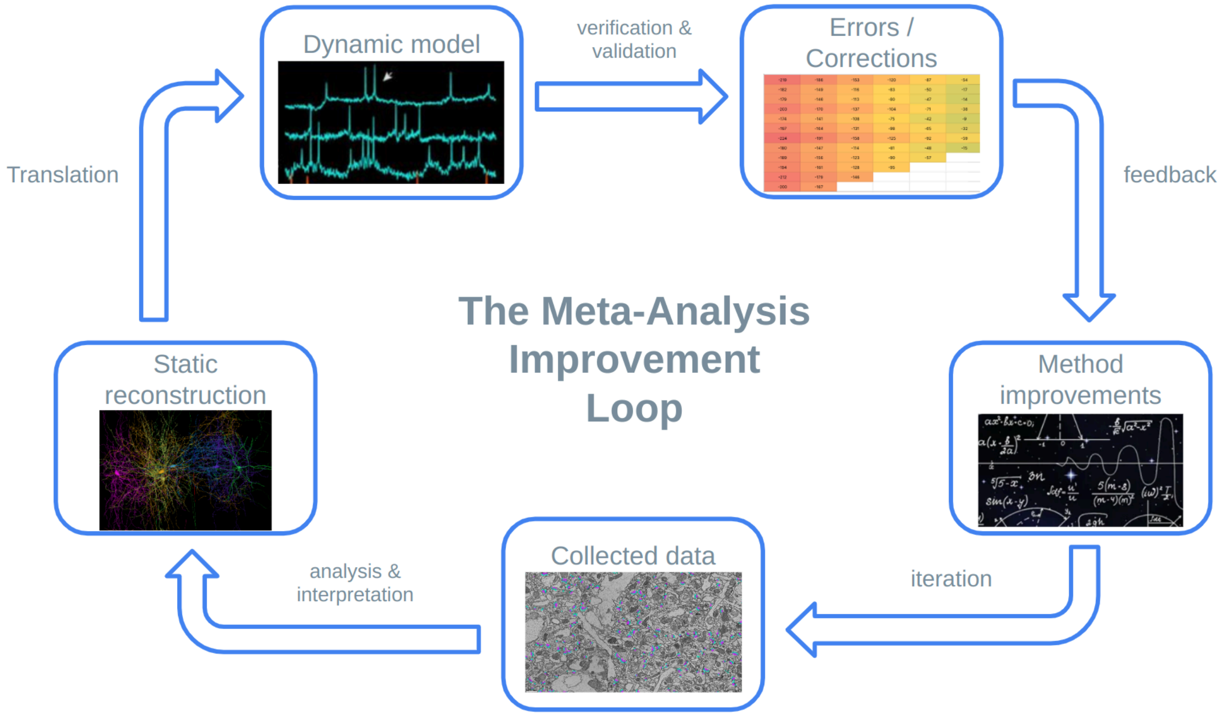

Figure 3: The meta-analysis improvement loop.

From connectome reservoir to example system

The neuronal outgrowth simulator, Netmorph (Koene et al., 2009), is configured to produce a local brain architecture (layers, cell types) for which neuronal outgrowth is then simulated until the detailed morphology of resulting neurons and their networks forms a reservoir of potential connections that is sufficient to express a cognitively meaningful neural circuit function (Fig.4). Pruning and tuning are then applied by presenting stimuli to the connectome reservoir, so that it does indeed express a desired neural circuit function, for example, associative memory with attractor dynamics using excitatory and inhibitory recurrent connections. This subsystem is hidden within an encompassing nonlinear dynamic system, i.e. neuronal circuitry not directly related to the function. All of this is embedded in a simulation of passive naturalistic components that do not directly contribute, but that may affect signals, e.g. extracellular materials.

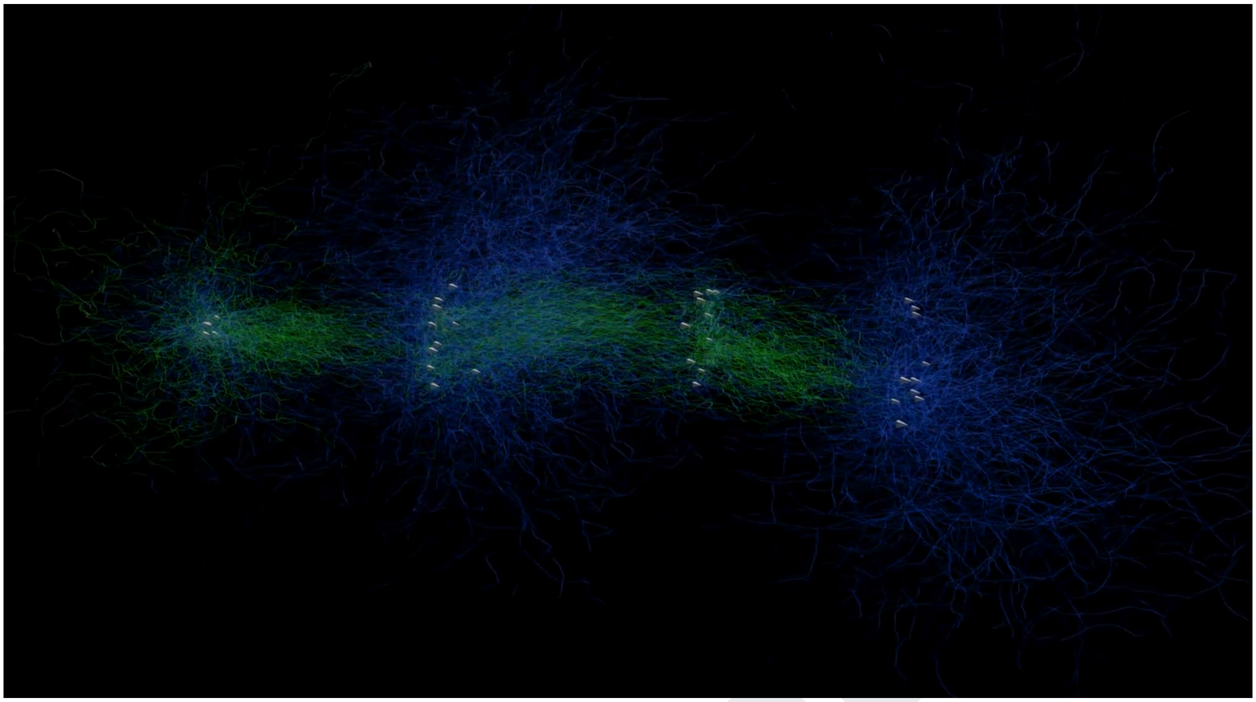

Figure 4: Connectome reservoir generated by Netmoph, including layers of Pyramidal neurons and interneurons.

Evaluating analyzed methods from data to emulation

Virtual data collection is carried out in the example system. This generates synthetic data for electron microscopy, calcium imaging, and electrophysiology. These data stacks are provided for analysis - without knowledge of the underlying ground-truth of the example system or its embedded meaningful function. The analysis methods being tested produce a proposed resulting emulation. The quality of the emulation is validated using direct knowledge of the example ground-truth, comparing the structure that was discovered, as well as the correctness of function. Functional testing can use the same I/O data sets initially used to tune the example system, as well as random stimuli for comparative sensitivity analysis. The evaluation points out strengths and weaknesses in the performance of analysis methods, which can be scored and compared across a wide range of approaches.

It is entirely possible to use the same protocol with randomly generated example systems. Our focus on evaluating analysis methods for their ability to correctly discover a meaningful subsystem is an explicit effort to steer the evaluation not towards optimizing performance on the reconstruction and reverse engineering of component-wise details, but instead towards goal-directed success criteria that are intended to emphasize cognitive experience, and to discover points of scale separation most relevant to those.

Generative Meta-Analysis is a development tool that helps pre-test analysis methods more easily and more rigorously before the best-performing methods are applied to biological data, where the ground-truth is rarely fully understood. Iteratively applying the protocol should lead to:

- Improved analysis methods, such as are required for functionalization.
- An improved understanding of the success criteria for brain emulation.
- An improved understanding of scale separation in the brain's neuronal circuits and mechanisms.
- Improved synthetic example generation for further meta-analysis as the insights are fed back into updates of the platform itself.

References

Arnaudon, A., Reva, M., Zbili, M., Markram, H., Van Geit, W., & Kanari, L. (2023). Controlling morpho-electrophysiological variability of neurons with detailed biophysical models. *iScience*, *26*(11), 108222. https://doi.org/10.1016/j.isci.2023.108222

Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Sciences*, *7*(4), 166–172. https://doi.org/10.1016/S1364-6613(03)00056-1

Belavkin, R. V. (2018). *Relation between the Kantorovich-Wasserstein metric and the Kullback-Leibler divergence* (Vol. 252, pp. 363–373). https://doi.org/10.1007/978-3-319-97798-0_15

Blackmore, S. (2017). *Consciousness: A Very Short Introduction*. Oxford University Press.

Chung, S. H., Sun, L., & Gabel, C. V. (2013). In vivo neuronal calcium imaging in C. elegans. *Journal of Visualized Experiments: JoVE*, *74*, 50357. https://doi.org/10.3791/50357

D'Esposito, M., & Postle, B. R. (2015). THE COGNITIVE NEUROSCIENCE OF WORKING MEMORY. *Annual Review of Psychology*, *66*, 115–142. https://doi.org/10.1146/annurev-psych-010814-015031

Dürst, C. D., Wiegert, J. S., Schulze, C., Helassa, N., Török, K., & Oertner, T. G. (2022). Vesicular release probability sets the strength of individual Schaffer collateral synapses. *Nature Communications*, *13*, 6126. https://doi.org/10.1038/s41467-022-33565-6

Gouwens, N. W., & Wilson, R. I. (2009). Signal Propagation in Drosophila Central Neurons. *Journal of Neuroscience*, *29*(19), 6239–6249. https://doi.org/10.1523/JNEUROSCI.0764-09.2009

Ikeda, K., & Bekkers, J. M. (2009). Counting the number of releasable synaptic vesicles in a presynaptic terminal. *Proceedings of the National Academy of Sciences*, *106*(8), 2945–2950. https://doi.org/10.1073/pnas.0811017106

Imbrosci, B., Schmitz, D., & Orlando, M. (2022). Automated Detection and Localization of Synaptic Vesicles in Electron Microscopy Images. *eNeuro*, *9*(1). https://doi.org/10.1523/ENEURO.0400-20.2021

Juavinett, A. L., Bekheet, G., & Churchland, A. K. (2019). Chronically implanted Neuropixels probes enable high-yield recordings in freely moving mice. *eLife*, *8*, e47188. https://doi.org/10.7554/eLife.47188

Ji, X., Elmoznino, E., Deane, G., Constant, A., Dumas, G., Lajoie, G., Simon, J., & Bengio, Y. (2024). Sources of richness and ineffability for phenomenally conscious states. *Neuroscience of Consciousness*, *2024*(1), niae001. https://doi.org/10.1093/nc/niae001

Kettunen, P. (2020). Calcium Imaging in the Zebrafish. *Advances in Experimental Medicine and Biology*, *1131*, 901–942. https://doi.org/10.1007/978-3-030-12457-1_36

Koene, R. A., & Hasselmo, M. E. (2007a). A reversing buffer mechanism that enables instances of retrospective activity in hippocampal regions CA3 and CA1. *2007 International Joint Conference on Neural Networks*, 1401–1406. https://doi.org/10.1109/IJCNN.2007.4371163

Koene, R. A., & Hasselmo, M. E. (2007b). First-In–First-Out Item Replacement in a Model of Short-Term Memory Based on Persistent Spiking. *Cerebral Cortex*, *17*(8), 1766–1781. https://doi.org/10.1093/cercor/bhl088

Koene, R. A., Tijms, B., van Hees, P., Postma, F., de Ridder, A., Ramakers, G. J. A., van Pelt, J., & van Ooyen, A. (2009). NETMORPH: A framework for the stochastic generation of large scale neuronal networks with realistic neuron morphologies. *Neuroinformatics*, *7*(3), 195–210. https://doi.org/10.1007/s12021-009-9052-3

Lappalainen, J. K., Tschopp, F. D., Prakhya, S., McGill, M., Nern, A., Shinomiya, K., Takemura, S., Gruntman, E., Macke, J. H., & Turaga, S. C. (2024). Connectome-constrained networks predict neural activity across the fly visual system. *Nature*, 1–9. https://doi.org/10.1038/s41586-024-07939-3

Metzinger, T. (2004). *Being No One*. MIT Press. https://mitpress.mit.edu/9780262633086/being-no-one/

MICrONS Consortium, Bae, J. A., Baptiste, M., Bishop, C. A., Bodor, A. L., Brittain, D., Buchanan, J., Bumbarger, D. J., Castro, M. A., Celii, B., Cobos, E., Collman, F., Costa, N. M. da, Dorkenwald, S., Elabbady, L., Fahey, P. G., Fliss, T., Froudarakis, E., Gager, J., … Zhang, C. (2023). *Functional connectomics spanning multiple areas of mouse visual cortex* (p. 2021.07.28.454025). bioRxiv. https://doi.org/10.1101/2021.07.28.454025

Otopalik, A. G., Pipkin, J., & Marder, E. (2019). Neuronal morphologies built for reliable physiology in a rhythmic motor circuit. *eLife*, *8*, e41728. https://doi.org/10.7554/eLife.41728

Parfit, D. (1984). *Reasons and Persons*. Clarendon Press.

Rall, W. (1959). Branching dendritic trees and motoneuron membrane resistivity. *Experimental Neurology*, *1*(5), 491–527. https://doi.org/10.1016/0014-4886(59)90046-9

Shapson-Coe, A., Januszewski, M., Berger, D. R., Pope, A., Wu, Y., Blakely, T., Schalek, R. L., Li, P. H., Wang, S., Maitin-Shepard, J., Karlupia, N., Dorkenwald, S., Sjostedt, E., Leavitt, L., Lee, D., Troidl, J., Collman, F., Bailey, L., Fitzmaurice, A., … Lichtman, J. W. (2024). A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science (New York, N.Y.)*, *384*(6696), eadk4858. https://doi.org/10.1126/science.adk4858

Shiu, P. K., Sterne, G. R., Spiller, N., Franconville, R., Sandoval, A., Zhou, J., Simha, N., Kang, C. H., Yu, S., Kim, J. S., Dorkenwald, S., Matsliah, A., Schlegel, P., Yu, S., McKellar, C. E., Sterling, A., Costa, M., Eichler, K., Bates, A. S., … Scott, K. (2024). A Drosophila computational brain model reveals sensorimotor processing. *Nature*, *634*(8032), 210–219. https://doi.org/10.1038/s41586-024-07763-9

Storm, J. F., Klink, P. C., Aru, J., Senn, W., Goebel, R., Pigorini, A., Avanzini, P., Vanduffel, W., Roelfsema, P. R., Massimini, M., Larkum, M. E., & Pennartz, C. M. A. (2024). An integrative, multiscale view on neural theories of consciousness. *Neuron*, *112*(10), 1531–1552. https://doi.org/10.1016/j.neuron.2024.02.004

Tavakoli, M. R., Lyudchik, J., Januszewski, M., Vistunou, V., Agudelo, N., Vorlaufer, J., Sommer, C., Kreuzinger, C., Oliveira, B., Cenameri, A., Novarino, G., Jain, V., & Danzl, J. (2024). *Light-microscopy based dense connectomic reconstruction of mammalian brain tissue* (p. 2024.03.01.582884). bioRxiv. https://doi.org/10.1101/2024.03.01.582884

Tuthill, J. C. (2009). Lessons from a Compartmental Model of a Drosophila Neuron. *The Journal of Neuroscience*, *29*(39), 12033–12034. https://doi.org/10.1523/JNEUROSCI.3348-09.2009

Yamada, R., & Kuba, H. (2021). Dendritic synapse geometry optimizes binaural computation in a sound localization circuit. *Science Advances*, *7*(48), eabh0024. https://doi.org/10.1126/sciadv.abh0024

Zhang, Y., Rózsa, M., Liang, Y., Bushey, D., Wei, Z., Zheng, J., Reep, D., Broussard, G. J., Tsang, A., Tsegaye, G., Narayan, S., Obara, C. J., Lim, J.-X., Patel, R., Zhang, R., Ahrens, M. B., Turner, G. C., Wang, S. S.-H., Korff, W. L., … Looger, L. L. (2023). Fast and sensitive GCaMP calcium indicators for imaging neural populations. *Nature*, *615*(7954), 884–891. https://doi.org/10.1038/s41586-023-05828-9

Zheng, Z., Lauritzen, J. S., Perlman, E., Robinson, C. G., Nichols, M., Milkie, D., Torrens, O., Price, J., Fisher, C. B., Sharifi, N., Calle-Schuler, S. A., Kmecova, L., Ali, I. J., Karsh, B., Trautman, E. T., Bogovic, J. A., Hanslovsky, P., Jefferis, G. S. X. E., Kazhdan, M., … Bock, D. D. (2018). A Complete Electron Microscopy Volume of the Brain of Adult Drosophila melanogaster. *Cell*, *174*(3), 730-743.e22. https://doi.org/10.1016/j.cell.2018.06.019