Journal of
Ethics and
Emerging Technologies

IEET

Article

Functional Tests Guide Complex "Fidelity" Tradeoffs in Whole-Brain Emulation

Charl Linssen [1,*] and Randal Koene [1]

[1] Carboncopies Foundation for Substrate-Independent Minds, 2108 N St #6380, Sacramento, CA 95816, USA
* Correspondence: clinssen@carboncopies.org

Abstract: The human brain can be understood as a vast network of neurons connected via synapses, the state of which is characterized by ion concentrations, phosphorylation patterns, receptor densities, etc. It is plausible that a mechanistic simulation at the scale of the whole brain (a "whole-brain emulation" or WBE) will be made, raising questions about moral status and personal agency. Creating a dynamical model of the brain presents a complex tradeoff between better performance, and data collection and operating costs. To make informed scientific, engineering, as well as personal decisions on the tradeoffs involved, a set of tests should be defined that quantify not only the dynamics of the underlying spiking neural network model, but also the performance of the emulated individual on a comprehensive repertoire of skills in a variety of domains. "Fidelity" can then be defined as a measure of how well the behavior of the model corresponds with the behavior of the original individual, or with respect to stereotyped brains. Models can subsequently be optimized to obtain the highest fidelity within the constraints of a given cost budget. However, an overall measure of fidelity is the outcome of a complex, high-dimensional optimization problem (that of choosing the parameters for a WBE) and remains in and of itself (as a measure or index) challenging to define. Different people and organizations are expected to make different tradeoffs based on a diverse set of criteria. Consequently, there can be multiple variants on offer for the translation from an original, biological brain to a WBE. If some variants are deemed cognitively superior, but are available only at a high cost, then this could have undesired socioeconomic effects where only those who are wealthy can afford the higher-tier emulations. However, competition between different WBE platform operators, which attempt to achieve the highest fidelity at the lowest cost, could help drive overall costs down. A framework of legal and ethical standards pertaining to model fidelity should be defined, which should recommend a minimum set of standardized tests.

Keywords: brain model; verification and validation; computational neuroscience, whole-brain emulation (WBE); substrate-independent mind (SIM); mind uploading

Citation: Linssen, Charl and Koene, Randel. 2025. Functional tests to guide complex "fidelity" tradeoffs in Whole-Brain Emulation. *Journal of Ethics and Emerging Technologies* 35: 1. https://doi.org/10.55613/jeet.v35i1.152

Received: 11/11/2024
Accepted: 8/03/2025
Published: 17/03/2025

Publisher's Note: IEET stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1. Introduction

The human brain is a physical object, the state of which varies over time. It can be measured, modeled by theory, and formulated mathematically as a dynamical system, allowing a computer to simulate its activity. As always in modeling, the model is a simplification of reality, typically involving certain approximations and assumptions. Modeling the brain is a difficult, multiscale challenge: there is an astronomically large number of neurons and synapses in the brain, and the spatial scales involved range from molecular interactions to electric fields at the scale of the whole brain, while timescales range from sub-millisecond processes, to learning that takes place over days to years. There is, however, no fundamental reason why a computer simulation of an individual human brain could not be built, that integrates so much of the biophysical details, that it would behave indistinguishably from the biological original. On a biophysical level, it would replicate the input-output behavior of single cells; by connecting these cells in a vast network on the basis of microscopy data, the activity of the network would verifiably (in a scientifically sound and falsifiable way) match that of the biological brain. When this biophysical simulation runs, the simulated brain would experience sensations, and exhibit cognition, self-awareness, a sense of identity, and agency via simulated or robotically mediated interactions with an environment through a certain embodiment ("avatar"). Potentially, this simulated individual would have a claim to human rights, and in general, would be a person just like you or I—except that instead of a biological brain, their brain function would depend on the electronic substrate of a computer. Analogous to a prosthesis, which can take over the function of a biological limb or an organ, a computer program becomes a neuroprosthetic of the original, biological brain. The brain data is obtained from a particular, individual human brain, so it is implied that the identity of that individual person is preserved by means of the computational neuroprosthetic. Furthermore, the person would experience a continuity of existence (such as by retaining their unique memories) from the time leading up to the scan to the point where the resulting WBE is operating (Sandberg and Bostrom 2008).

The term "simulation" refers to mimicking the evolution through time of some measurable quantity of the physical system (the brain) on a computer. Instead of a potentially superficial facsimile that only reproduces a select number of outwardly measurable quantities, an *emulation* is a simulation that is based on mimicking the internal mechanisms of the process (Sandberg and Bostrom 2008). In the case of a brain emulation, this means going beyond a statistics-based model (Rothblatt 2013) by replicating aspects of the structure and function of a brain, such as the architecture of a network of spiking neurons, that allow system identification and reconstruction even for system behaviors not seen in a limited dataset. A mechanistic model on the scale of the entire human brain is referred to as a "whole-brain emulation" (WBE; Koene 2012).

The most foreseeable path to WBE involves high-resolution microscopy of brain tissue, followed by computer-aided segmentation of the scans and reconstruction into a three-dimensional volume (Xu et al. 2020; Dorkenwald et al. 2023; Collins et al. 2025). The features typically observed in these scans are at the scale of single cells and synapses: individual neurons and their ramifications can be readily distinguished. To further develop the resulting anatomical model into a dynamical model that can be simulated on a computer, functional properties of the cells have to be known, perhaps derived from known structure-function relationships (Allen 2022) or additional types of scans, such as immunohistochemistry (Fulton and Briggman 2021). Parameter fitting and optimization may need to be carried out to allow the digital network model to perform computational tasks (Lappalainen et al. 2024) and do so in a way emulating the biophysical mechanisms of the real brain (Shiu et al. 2024). The overall process of scanning and digitisation of a brain, and elaborating it into a dynamical model that is then simulated on a computer, is popularly known as "uploading".

We note that there are rich, ongoing debates in the literature regarding the feasibility of WBE (see for instance Chalmers 2014), in particular regarding the technical feasibility and questions of consciousness and preservation of personal identity (Minerva 2021). It is outside the scope of this paper to address and provide counterarguments for these critiques; for the purposes of this paper, the theoretical and practical feasibility, and the eventual realization of WBE, will be postulated to be an accepted truth (an axiom), in the interest of serving as the basis for further reasoning and discussion.

2. Tradeoffs in design decisions

Computational models must be formulated using a precise syntax and semantics and in a way amenable to computer simulation. First, a sensible set of features has to be selected to represent the system. The rules of evolution in time of these features have to be established, and the values of system parameters need to be derived from measurements or be otherwise inferred. Then, the model (features and rules of evolution) must be implemented in a form that can be run on a computer. In each of these steps, there is a complexity tradeoff: incorporating more parameters and state variables into the model might improve its accuracy and predictive value, but requires more data to be gathered in order to assign numerical values to parameters and to elucidate interrelationships between state variables. The time taken to achieve a good model fit to the empirical data, a constrained optimization problem, increases drastically with the number of interacting variables (a combinatorial explosion). Storage requirements increase with the number of variables and parameters of the model. When running the emulation, the state update dynamics of the model can be computed with greater numerical precision, but at the cost of requiring additional computational power or wall-clock time.

Although in the long-term future, the cost of computation may tend towards zero (Bostrom 2003), and similar arguments can be made for the mass, volume and energy requirements of the computational substrate, the tradeoff argument will continue to apply as long as there is at least some non-zero cost to computation (as well as the requisite storage, communication of data across a network, etc.) This will certainly remain the case in the foreseeable future, especially until the point where novel computational paradigms such as reversible and quantum computing could be adopted, although their usability for brain simulations has yet to be demonstrated. Even in a distant future, a constant striving for improvement, and contention over finite resources could be sufficient to drive an evolutionary process that involves resource/performance tradeoffs (see discussion about the "Red Queen hypothesis" in the Discussion section).

Thus, in practice, a complex tradeoff must be made involving many options and possibilities. We illustrate this complexity by providing several concrete examples of tradeoffs related to model construction and model simulation, respectively, in the following two subsections.

2.1. Design decisions related to model construction

Computational models are representations of real physical systems, with which quantitative, testable predictions about those systems can be made (Horsman et al. 2014). Models are constructed based on a set of criteria that suit the context of the application or purpose of the model, such as phenomena or behaviors the model aims to predict. The model may or may not make useful predictions beyond that context, but it is not guaranteed by its construction; a model is an idealization and is intended to operate well within a given application range.

A real-world system that is being modeled thus does not have a prescribed unique or inherent scale, formulation, or level of description (as already observed in Rosenblueth and Wiener 1945). However, the emulations of interest to us should include explicit representations of underlying mechanisms and their interactions, through a system architecture that corresponds with the real-world system. They should not reproduce only externally observable behavior, but operate on the basis of similar principles and laws of physics, that describe how the state of the system changes across time. For instance, a weather prediction model could be based on a simple statistical model of temperature and precipitation on previous days, or it could involve a detailed simulation of atmospheric currents, although both might make similar predictions. If a model has a mechanistic basis, then it allows us to speak about a *level of detail:* more aspects or components of the system's mechanism can be incorporated into the model to provide a greater level of detail. These components are compatible and mutually interact by virtue of being analogues of their real-world counterparts. Mechanistic models with a higher level of detail would be expected to have greater predictive value (in this example, to be better predictors of the upcoming weather) and to generalize better to situations that they were not explicitly trained or adjusted for.

Dynamical models of the brain can also be defined at various levels of detail; the 2008 Roadmap to Whole-Brain Emulation considered eleven resolutions (Sandberg and Bostrom 2008), including tracking individual molecules, such as neurotransmitters at the highest resolution, and treating an entire population of neurons as a single partial differential equation at the lowest resolution.[1] Several ways in which model resolution may affect the emulation are suggested by existing insights into brain physiology, current practices of modeling and simulating, and known links with function and behavior. We will review several concrete examples of such insights in the next paragraphs, before coming to a general conclusion.

A first example comes from the volumetric reconstruction of all neurons and synapses in a *Drosophila* larva brain, where manual proofreading suggests errors of omission in the reconstruction of synaptic connectivity (Winding et al. 2023). In general, false positives and false negatives may be rare in proportion to the volume imaged, but depending on their distribution and on the robustness of information representation in neural networks, such errors may be cumulatively significant. As synaptic connectivity is posited to be the physical mechanism for the retention of memories (Frankland et al. 2019), errors of omission are a concrete example where specific details of a data analysis step have direct impact on an individual's personal identity, as it could erase or distort existing memories, presenting as retrograde amnesia. This implies that there is a tradeoff to be made: how much compute time and other resources do we dedicate to the 3D reconstruction algorithm, especially when running the algorithm for longer and longer times runs into diminishing gains in reconstruction accuracy?

In the compartmental modeling method, which is highly successful in predicting biological firing activity (Herz et al. 2006), realistically modeling biological features of real neurons is the principal aim. Typically, compartmental neuron models are based on reconstructions from (electron) microscopic scans of individual cells, and automatically processed into thousands of discrete compartments (Bates et al. 2020). However, large numbers of compartments means increased computational cost. The models can typically be simplified to contain fewer compartments, while retaining electrical-circuit characteristics and functional responses to inputs (Wybo et al. 2021; Amsalem et al. 2020). The simplification process is gradual, with error measures (compared to the baseline model with the largest number of compartments) gradually increasing with decreasing compartment count. There are no well-established rules for what compartment count is adequate, but the count is defined depending on other factors such as the conditions under which the simplification strategy is evaluated (for instance, response to input spike patterns of varying complexity).

Brain function relies on highly dynamic, activity-dependent processes that switch genes on and off. These can lead to profound structural and functional changes, involving the formation of new, and elimination of unused synapses, changes in cytoskeleton, receptor mobility as well as energy metabolism. Cognitive ability may crucially depend on how efficiently neurons can regulate these processes (Goriounova and Mansvelder 2019). It was found that a genetic overall factor explained, on average, over 58% of the variance across seven cognitive tests in people with European ancestry. Genome-wide association studies (GWASes) were conducted on intelligence test scores to identify

1 Let us assume that at a given level of description, the constituent components can be described by on the order of 10-1000 state variables; it does not make sense to talk about a cell-level model that contains so many state variables that it becomes essentially equivalent to a more detailed (for example, proteome-level) description.

associated genetic loci. Using large sample sizes, of over 200,000 participants, hundreds of genetic loci significantly associated with intelligence were found, each of which accounts for a tiny proportion of overall intelligence (Deary et al. 2021). For some of these genes, the mechanisms by which they influence cognition are known. For example, a common functional polymorphism in the *COMT* gene has been associated with improvements in working memory performance. Catechol-O-Methyltransferase, the protein coded for by the gene, breaks down dopamine in the prefrontal cortex, whereas increased dopamine availability is associated with better memory performance. Similar polymorphisms exist for other genes that are known to encode for proteins involved in cellular maintenance and repair, such as Apolipoprotein E (*APOE;* Dongés et al. 2012). These findings match what is known from research into the consolidation of memories, which indeed suggests a complex process between storage and retrieval, relying on the efficacy of protein synthesis and turnover, suggesting involvement of an intricate intracellular signaling network (Nader et al. 2000).

In general, intracellular pathways of second messengers and downstream effects (Sorokina et al. 2021), phosphorylation of various proteins, dendritic spikes, cellular morphology and microdomains, the stochastic properties of ion channels (Cannon et al. 2010), DNA translation and epigenetic regulation, glia, the dozens of known neurotransmitters and neuromodulators (Sandberg and Bostrom 2008), retrograde signaling (Citri and Malenka 2008), and so on, constitute a vast swath of intra- and extracellular mechanisms that are implicated in the computational function of neurons and synaptic plasticity. An emulation that incorporates a higher level of detail and increasing number of these mechanisms would demonstrably increase the similarity to the biological original, and have greater explanatory and predictive value as a model, but requires additional computational resources. Some biophysical processes, such as those related to cell metabolism, or the regulation of local oxygenated blood flow in the neuropil, could have effects that are very hard to detect on a functional (behavioral) level without specific testing, or not deemed sufficiently relevant to neural computation on the basis of theoretical considerations (or may even be undesirable, such as finite blood oxygenation).

2.2. Design decisions related to model simulation

The outcome of the design process in paragraph 2.1 is a formal description of a set of interacting mechanisms, that is, a model containing state variables, parameters, and a set of rules that describe how the state of the system changes across time. Simulation software takes the model description as input, and computes the numerical outcomes of these rules, advancing the state of the simulation in (simulated) time, while also providing interfaces for reading out the state of the system and providing external stimuli, for example in a closed loop physics simulation of a simulated or real-world body avatar. Numerical accuracy of the simulation methods can have a large impact on network dynamics, and by extension, how (well) the network functions.

Many physical models of dynamical processes are formulated in terms of a set of ordinary differential equations (ODEs), which describe how quantities in the model change over time. Simulating these dynamical systems on a digital computer requires the discretization of time into discrete steps, and taking jumps from one step to the next; the state of the system is typically only well defined before and after each jump (Koene 2012). The size of the timestep Δt (or in other words, the temporal resolution of the simulation) is a free parameter chosen by the modeler and is often purely based on heuristics, or possibly, based on a maximum error bound that is itself heuristically chosen. A smaller timestep (up to a certain point) generally results in better accuracy, that is, a better approximation of the continuous-time ODEs by their discrete approximations, but as halving the timestep typically doubles the amount of computation required, Δt is a centrally important quantity to be chosen optimally.

On a computer, the variables describing the system have to be represented in a finite bit resolution. On modern computers, a default of 64 bit floating-point resolution (IEEE 2019) provides 11 bits for the exponent and 52 bits for the mantissa. Some models use only 8 bits of resolution (Müller and Indiveri 2015) or even fewer (down to 1 bit). Because of the very large number of synapses in the model, small savings in storage and computation per synapse translate into considerable gains on the network level, in terms of hardware requirements, processing efficiency and power consumption (Pfeil et al. 2012; Bartol et al. 2015). Novel analog neuromorphic hardware could impose similar constraints. Although the discussion on bit resolution and timesteps applies to digital and not analog computers, the latter have their own fundamental limits on representational accuracy in terms of electrical noise, component variability, unintentional non-linearities and parasitic effects such as electrical coupling between adjacent components (Petrovici et al. 2014).

Taken together, the degree to which the output of a simulation approximates the true solutions to the mathematical model upon which it is based, as well as the degree to which the mathematically ideal model approximates the real, biological system upon which it is based, are the outcome of a high-dimensional set of mutually interacting design decisions, many of which have a graded (as opposed to all-or-nothing) nature. It is not clear *a priori* which values are adequate for the model; for example, whether increasing the number of bits representing floating point numbers, or doubling the number of compartments in a model, should be considered to give equal contributions to improved model function, given the extra computational burden that is implied. Although effects can be subtle individually, even small effects could have large interactions when combined, as seen in the genetic studies. Some individuals may consider even subtle effects to be worth the investment—although there might be mainstream consensus, it would also be conceivable to operate under a different set of (personal) norms, in the same way that people have different standards about their diet and supplement intake.

3. Functional tests guide design decisions

In section 2, we argued that a complex, multidimensional set of tradeoffs exists in all stages of constructing a WBE, from the data gathering, to model construction and simulation methods. Due to the high dimensionality and complexity of this optimization problem (that is, to choose the optimal model and simulation parameters in the space of all possible design decisions), there could be many viable choices, and it is not clear which level of detail, or which variables (and which dynamics) are required in the WBE model. We now address how informed decisions can be made on these tradeoffs.

In general, models are created in the context of a set of target phenomena that the model is set out to capture; thus, the targets that are defined should be taken into careful consideration and can be a source of divergent design decisions during data collection and model construction. Having a set of targets implies that the degree to which the targets have been met by the model needs to be quantified. This raises the question of which tests will be used, and to what degree the actual outcome of a test is expected to meet a certain threshold or a set of criteria for the test to be passed. Tests could be defined on the basis of comparing simulated quantities to empirically measured values. This gives rise to the term *fidelity:* the degree of correspondence between characteristics of the real person, and the WBE that was created on the basis of their brain scans. General human characteristics, as well as specifically individual traits, skills and memories, are expected to be faithfully replicated by the emulation. To allow a broad basis for comparison, these characteristics can be quantified by profiling individual scores across a battery of standardized tests. In high fidelity simulations, the profile would be expected to show an overall high degree of correspondence between the simulated and empirical values.

3.1. Tests of the network dynamics

Functional testing of spiking neural networks can be performed on different spatial and temporal physiological scales. The link between the dynamics at these different scale levels might be fundamentally inaccessible to inquiry; hence, complementary approaches are needed. There is no trivial way to compare the dynamics of high-dimensional systems. Instead, a battery of tests is required, each of which contributes to an overall "score" for a network. Some of the measures include: the overall level of firing activity of individual neurons or particular groups of neurons in the network; the distribution of interspike intervals; the self- and cross-correlation functions (pairwise and higher-order), which quantify correlations between spike counts of binned spike trains (on different analysis timescales); eigendecomposition of the network activity correlation matrix; methods for detecting particular, repeating spatiotemporal spike patterns; measures for quantifying synchronicity in spiking and phase locking; and tests for responses of the network to spontaneous network activity (Gutzen et al. 2018) as well as external stimuli of various kinds, measuring for instance the complexity in the resulting network response (Bayne et al. 2024). Running an emulation with too large a timestep can result in artifactual synchronization of neuronal spiking in the network, an undesired phenomenon which gradually disappears as the timestep becomes smaller. The value of the timestep should therefore be carefully considered (Hanuschkin et al. 2010).

3.2. Cognitive and behavioral testing

Further functional tests can be defined that specifically address design decisions related to model construction (section 2.1). Validation, the process of determining whether the chosen model is an adequate representation of the real-world system for its purpose (Gutzen et al. 2018), should be addressed through a series of standardized behavioral tests (Eliasmith and Trujillo 2024). Many types of standardized tests exist to assess cognitive performance, that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, and learn from experience. It is not merely book learning, or test-taking, which can be considered a narrow academic skill. Rather, it reflects a broader and deeper capability for comprehending our surroundings and performing rapid and adaptive problem solving, and should be understood more broadly to include skills in multiple domains (Deary et al. 2021).

In general, tests may score on the basis of observed behavior (like error rates on a task), self-report, and reaction times. Many formal, standardized assessments have been developed, covering a variety of different domains, such as for assessing a person's executive function and ability to concentrate on a task, sensori-motor function, visuo-spatial processing, motor skill, memory and learning, language ability, social intelligence, and creativity (Goldstein and Hersen 2000; Goldstein et al. 2003). We briefly discuss a few tests from these categories for illustration.

Mazes test spatial navigation tasks and visual processing. Visual acuity tests were used to compare the performance of a fully-detailed, and a simplified version of a large-scale primary visual cortex model (Bileh et al. 2020). Other geometrical tasks include, for example, skill at performing mental rotations of a three-dimensional object. Perception-oriented tasks should be complemented by dexterity and motor proficiency tasks, for instance, from simple manipulations like "peg board" tests, to playing a musical instrument or typing on a computer keyboard at a certain level of proficiency.

Working and long-term memory can be tested by means of Delayed Matching-to-Sample (DMS) and *N*-back tests. In DMS, participants are required to match a sample stimulus to a target stimulus after a delay period, for example in the Benton Visual Retention Test, which involves visual memory of an abstract figure. The difficulty of the task can be manipulated by varying the delay period or the number of distractor stimuli. In the *N*-back task, participants monitor a sequence of stimuli (for example, letters or numbers) and respond when the current stimulus matches the one presented *N* steps back in the sequence. The difficulty can be manipulated by varying the value of *N*. Similar tests exist in the auditory domain, such as for verbal and nonverbal auditory memory.

Natural language skills can be tested for by antonym and synonym vocabulary, question answering, sentiment analysis, and textual entailment (Wang et al. 2019). Language tests assess both the structure as well as content of language production, including fluency, articulation, prosody, grammar, word finding, proportion of informative words produced, in tasks such as open-ended question answering, story comprehension and retelling, and picture description (Goldstein et al. 2003).

The Wisconsin Card Sorting Test measures flexibility in cognitive strategies in response to changing patterns of feedback. Tower of London (ToL) is a test that assesses planning and problem-solving skills. Behavior Rating Inventory of Executive Function is a questionnaire that measures behavioral regulation, organization and metacognition. Further tests address directing of attention and cognitive control, such as the Stroop test, in which participants are presented with color words (say, "red") printed in congruent or incongruent ink colors (in this example, say, either red or green). Participants must then identify the ink color while ignoring the word's meaning.

Further tests exist for aspects of social cognition, theory of mind and prosocial behavior. Personalized tests can be used, for example, testing retrieval of specific, personal memories. Aptitude tests can be administered for the specific field(s) that the individual is an expert in. In the field of consciousness studies, several functional tests are defined, including the detection of statistically second-order outliers and appropriate higher-level understanding of movie segments (Bayne et al. 2024). Tests can even be administered of susceptibility and responses to mystical experiences, which can be induced through (simulated) pharmacology, but also through meditation and ritual. These include the Mystical Experience Questionnaire (Barrett et al. 2015) and Cole's Spiritual Transformation Scale (Cole et al. 2008).

Several types of test require a combination of skills to perform, for example, the Paced Auditory Serial Addition Test (PASAT) asks subjects to add a series of numbers together and keep the running total in memory, which requires complementary skills in auditory processing, executive function, working memory, and mathematical skill. The Rey–Osterrieth Complex Figure (ROCF) test asks subjects to replicate a complex figure drawing from memory, which relies on skills in visual perception as well as motor control (Goldstein et al. 2003).

Note that all test metrics are intended to inform a correspondence between the model and the original, rather than being itself a score to improve. For instance, if the individual was afflicted with dyslexia or dyscalculia, then their WBE will be expected to be similarly affected (but see the discussion about cognitive enhancements in the Discussion section below).

3.3. Achieving consensus on test batteries

In general, a picture emerges that a large and diverse array of test approaches exist. Some of the tests in sections 3.1 and 3.2 have been discussed in the peer-reviewed literature for several decades, having been widely used and thoroughly validated. In spite of this, it is not clear which of these tests are strictly necessary, which are overlapping in what they test for, or which further tests may need to be added for a WBE. When tests are formulated in terms of spikes, it can be hard to interpret what they test for, as spiking data can not yet in general be "decoded" (the semantic meaning read out and interpreted). Furthermore, many test outcomes are formulated in terms of one or more numerical scores, and it is not clear which scores should be deemed adequate.

Tests should be checked for *reliability* in outcomes. Reliability in a test is the extent that its results are consistent, irrespective of time of administration, proctor, test form, etc. Although individuals taking tests may have (unconscious) incentives or inclinations for faking good or bad, exaggerating, defensiveness, carelessness, or responding randomly, tests can be designed to measure inconsistencies and exaggeration in the responses. Several quantitative methods exist for evaluating the reliability of a test, for instance, splitting the test in two even halves and comparing results on each half. Sensitivity analyses or comparison with surrogate (counterfactual or alternate) models can help to provide a reference point. These tests would be especially useful for checking design decisions related to the simulation method (section 2.2), as surrogate models with different parameters (bit resolution, timestep, etc.) can be easily run and compared.

Tests should additionally be checked for *validity*, which is the extent that a test measures what it purports to measure, including being complete and representative for the testing domain addressed by the test, and that the test predicts or correlates well with real-world outcomes. When using a battery of tests to assess characteristics in one and the same domain, tests that purport to measure the same thing should have a high correlation in outcome (and vice versa). This can be used to improve the testing procedures themselves and help to establish a mutually consistent set of tests; even across network levels (section 3.1) and behavioral levels (section 3.2).

As there are many detailed individual design decisions to be made, a multidimensional, complementary system of tests is suggested. However, there is no universal, canonical battery of tests, neither on the network dynamics level nor on the behavioral level. The best practices of testing of neural network dynamics (section 3.1) and testing on the individual, behavioral level (section 3.2), continue to develop and evolve. New scientific and anthropological insights should continue to be integrated into the testing battery. For instance, only relatively recently, recommendations were formulated to use specialized techniques for assessing persons who are culturally distinct from the populations on which psychological and neuropsychological tests have been normed, and for whom these norms are likely to be misleading (Goldstein et al. 2003).

Determining an appropriate testing framework thus itself presents a complex set of choices, where a gradual consensus emerges only through the scientific process, as well as societal debates and the moral norms and standards at a given time, such that they remain subject to continuous change and development as new insights emerge. Decisions about the appropriate testing framework to use are furthermore influenced by personal epistemological choices about, for example, what information sources to trust. Thus, even a comprehensive battery of tests should be considered a necessary, but not sufficient criterion for assessing the fidelity of a WBE.

Furthermore, major changes in testing outcomes would be expected to stem as a natural consequence from the radical existential change that comes from being uploaded. Personality, needs, desires and goals, could all reasonably be expected to undergo a shift on the basis of moving from leaving a biological body behind to become uploaded. This shift will confound the basis of comparison, especially when it comes to personality or spirituality tests, and could be an argument for taking an initial test battery directly after the upload procedure and under appropriately defined conditions, as it would then focus on the objective changes due to the new substrate, while minimizing the influence on the subjective perception of the self and paradigm shifts that happen over time.

4. Discussion

Like a meteorological model of the atmosphere, a model of the human brain can never include all the details of the thing that it describes, but it can contain enough detail to be useful. Utility can be expressed as a set of design targets or phenomena that the model is set out to describe, and a corresponding set of tests that demonstrate that the targets have been met by the model (such as accurately predicting the weather). The set of targets should include those at the level of network dynamics (corresponding to section 3.1) as well as those at a behavioral level (corresponding to section 3.2).

Even after the design targets have been determined, a large number of choices remain to be made about which biophysical mechanisms should be included in the model (and with which parameters) in order to meet those targets. Given the synaptic and neuronal microstructure that is evident in high-resolution scans of the brain, and how well we can predict (and control) spiking at the single-cell level, the cellular level of description seems like a *prima facie* plausible level of description for a WBE. However, the brain is a highly complex dynamical system, with many processes that are relevant for its computational role, and many processes that are not—and sometimes, even this distinction can be hard to draw, like the finite blood oxygenation discussed in section 2.1. Increases in model complexity (typically involving an increased number of state variables and parameters) might be associated with an improvement in numerical and behavioral performance of the model, and improve the quantitative agreement between an emulation and the real system it aims to model—in our terminology, the "fidelity"—but at the same time, tradeoffs must be made, because improved fidelity is offset by the cost, difficulty, time investment, risk, energy, and effort involved with data collection, model fitting and operating the emulation substrate. These tradeoffs may be driven by cost considerations, be based on pragmatism, personally held beliefs, or be heuristic in nature, even when the purpose of the emulation is ostensibly the same.

To help individuals and organizations take informed decisions about model fidelity, development of a comprehensive battery of tests is suggested (section 3.1 and 3.2). Developing this system of tests is a complex task, and although a certain scientific consensus is expected to be achieved, a range of individual and organizational choices, made on a vast array of model implementation details and configuration parameters, will result in the emergence of a diverse spectrum of different approaches to WBE (as argued in section 3.3), even under the assumption (made in section 1) that WBE will generally be a widespread success.

Most fundamentally, there is no current agreement about what the moral status of the upload is. There is no agreement about which specific cognitive capacity (or its level of sophistication) is a necessary and sufficient condition to grant moral status. Suggestions include consciousness, sentience (reasoning), the presence of a concept of self and self-awareness, the capacity to engage in self-motivated activity, and the capacity to understand oneself as a continuing object of experiences (Minerva 2021). Formulation of an adequate set of tests for these capacities is clearly a task of grave importance if the test results are going to be used to determine a being's moral status.

Provided that moral status is granted to be equivalent to that of a non-uploaded person, a unique set of ethical considerations arises. For example, in circumstances where the person was suffering, they could ask their operator to end the simulation and destroy the backups, causing them to cease to exist, akin to euthanasia. However, different substrate operators might have different views on what constitutes an acceptable moral framework, and a set of regulations may be in place that governs these decisions, allowing some and restricting other options available to the uploaded individual (Minerva 2021). Clear liabilities should be established, for example in cases where defects are found in the model to start with as a result of difficulties encountered in the upload process, or that, due to long-term simulation in too low a fidelity, the model could incur damage or be otherwise disadvantaged, which the substrate operator could then be culpable for.

Relatedly, different substrate operators might provide different tiers of making back-ups of the WBE model: in case a calamity strikes the data center hosting the computational hardware that is running the WBE model, an earlier, backed-up version of the model could be reinstated (Linssen and Lemmens 2016). Depending on the tier, back-ups could be made at different intervals or with further-reaching safety guarantees. The integrity of WBEs could also come under direct attack: imagine a blackmail computer virus that infiltrates the platform hosting the WBE, and tortures the digital entity (inflicting physical pain on the simulated body, or causing mental distress), unless a ransom is paid. Just as we keep improving the security of our existing software, to better protect them from hacking attacks, at the same time hackers keep improving their skills and techniques to break into these systems. Possibly, a state of full (cyber)security can never be achieved, but is in a dynamic equilibrium with attackers, a phenomenon which is known from evolutionary biology as the "Red Queen hypothesis". Back-up arrangements and defense from computer viruses could constitute novel types of emulation tiers offered by substrate operators, closely analogous to tiers pertaining to fidelity.

The complex issues that arise related to privacy of WBEs are already being discussed at present, in debates about neurotechnologies that offer the possibility of directly decoding thoughts from neural activity (Baselga-Garriga et al. 2022). For a WBE, the neural activity of the whole brain would be easily accessible and used in advanced decoding tools. It would thus be comparatively easy for an external entity to access the content of memories, experiences and mental states of the uploaded brain (Minerva 2021). Furthermore, if the individual was instantiated in a (robotic) avatar through a wireless link, the body as well as the link create additional privacy concerns (Minerva 2021) as well as a new set of fidelity-versus-cost tradeoffs pertaining to the avatar and the wireless link.

On a free market, multiple operators and providers of WBE substrate infrastructure (hardware and software) could exist, offering different tiers or packages, that offer WBEs with a diverse set of options at different prices. In addition to qualitative differences in model architecture, such as which biophysical mechanisms are included or not, one of the most readily apparent WBE parameters to compete on is acceleration of the emulation with respect to real time. Provided that enough computing power is available, the emulation could almost trivially be run any number of times as fast or as slow as "real" (wall-clock) time (Hanson 2016). Migrating a WBE model from one operator to another would in general be possible, depending on whether the WBE model architecture is comparable between the two operators. If model architecture is tightly coupled to the specific (neuromorphic) emulation hardware used, then difficulties could arise and give rise to "operator lock-in". Additionally, in a free market, monopolies of WBE operators could ensue, which would be undesirable because it reduces consumer choice and technological innovation, and results in higher prices (Fagella 2020).

In this scenario, WBE could also become a driver of socio-economic inequality, as those who cannot afford a high-tier emulation fidelity are forced to live in a state of mental or physical impairment compared to those who can. Lower-tier emulations could predispose their subjects to mental disorders, or create new classes of mental disorder not found in nature, brought about by a particular set of design decisions in the model (Krzanowski and Krzanowski 2023). In contrast, high-tier emulations provide a competitive advantage, for example in business or trading (even if only by being able to respond faster due to being able to afford a more highly accelerated emulation), which would lead to an ever-growing disparity between high and low socio-economic status.

An appropriate, overarching legal and ethical framework needs to be established, extending the current set of "NeuroRights" (Baselga-Garriga et al. 2022), to address the concerns raised above. For instance, regulations could be put into place that grant the ownership of the hardware on which the brain is run to the emulated person, in the same way as we have ownership of our bodies. Regulations regarding minimum acceptable fidelity (as well as extra provisions such as virus protection and back-up policies) could constitute "human rights"-type laws for WBEs.

WBE creates a lot of freedom in general for further modifications to the model, some of which might demonstrably boost intelligence, akin to present-day nootropics or drugs. In addition to time acceleration, interventions might involve dynamically switching between different sets of configuration parameters of the model while it is running, as it is easy in a simulation to, for instance, replace all of the neurons of a particular cell type with another variant, picked from a library, that has less or more fidelity to the original, even on a brain region-by-brain region basis and on the basis of a changing set of demands over time. Offering a dynamic control over these parameters constitutes yet another dimension of cost-versus-fidelity tradeoff tiers. Which of these will be offered by a given platform operator is again an open question: some might consider certain options and possibilities unethical (like "wireheading", the act of directly stimulating the brain's reward center), whereas others could consider interventions like this ethical or even desirable, as a right to mental self-determination, and the expression of humankind's fundamental drive to decrease suffering and improve themselves and the world around them (Bublitz 2013; Pearce 2023).

Author Contributions: C.L, R.K: Conceptualization and writing. All authors have read and agreed to the published version of the manuscript.

Funding: N/A

Institutional Review Board Statement: N/A

Informed Consent Statement: N/A

Data Availability Statement: N/A

Acknowledgments: N/A

Conflicts of Interest: The authors declare no conflict of interest.

References

(Allen 2022) Allen Institute for Brain Science, Allen Institute for Cell Science. 2022. "Allen Brain Atlas: Data Portal." https://celltypes.brain-map.org/ (accessed on October 17th, 2022).

(Amsalem et al. 2020) Oren Amsalem, Guy Eyal, Noa Rogozinski, Michael Gevaert, Pramod Kumbhar, Felix Schürmann and Idan Segev. 2020. "An efficient analytical reduction of detailed nonlinear neuron models." Nat Commun 11, 288

(Ashton & Lee 2007) Ashton M.C and Lee K. 2007. "Empirical, theoretical, and practical advantages of the HEXACO model of personality structure." Pers Soc Psychol Rev. 2007 May; 11(2):150-66

(Barrett et al. 2015) Barrett FS, Johnson MW, Griffiths RR. 2015. "Validation of the revised Mystical Experience Questionnaire in experimental sessions with psilocybin." J Psychopharmacol. 29(11):1182-1190

(Bartol et al. 2015) Bartol T.M. Jr., Bromer C., Kinney J., Chirillo J.N., Bourne J.N., Harris K.M. and Sejnowski T.J. 2015. "Nanoconnectomic upper bound on the variability of synaptic plasticity." eLife 4:e10778.

(Baselga-Garriga et al. 2022) Baselga-Garriga, C., Rodriguez, P. and Yuste, R. 2022. ”Neuro Rights: A Human Rights Solution to Ethical Issues of Neurotechnologies.” In: López-Silva, P., Valera, L. (eds) ”Protecting the Mind. Ethics of Science and Technology Assessment.” Vol 49. Springer, Cham.

(Bates et al. 2020) Bates A.S., Schlegel P., Roberts R.J.V., Drummond N., Tamimi I.F.M., Turnbull R., Zhao X., Marin E.C., Popovici P.D., Dhawan S., Jamasb A., Javier A., Capdevila L.S.C., Li F., Rubin G.M., Waddell S., Bock D.D., Costa M. and Jefferis G.S.X.E. 2020. ”Complete Connectomic Reconstruction of Olfactory Projection Neurons in the Fly Brain.” Current Biology, Volume 30(16):3183-3199.e6.

(Bayne et al. 2024) Bayne T., Seth A.K., Massimini M., Shepherd J., Cleeremans A., Fleming S.M., Malach R., Mattingley J.B., Menon D.K., Owen A.M., Peters M.A.K., Razi A. and Mudrik L. 2024. ”Tests for consciousness in humans and beyond.” Trends in Cognitive Sciences

(Bileh et al. 2020) Yazan N. Billeh, Binghuang Cai, Sergey L. Gratiy, Kael Dai, Ramakrishnan Iyer, Nathan W. Gouwens, Reza Abbasi-Asl, Xiaoxuan Jia, Joshua H. Siegle, Shawn R. Olsen, Christof Koch, Stefan Mihalas and Anton Arkhipov. 2020. ”Systematic Integration of Structural and Functional Data into Multi-scale Models of Mouse Primary Visual Cortex.” Neuron, Volume 106, Issue 3, Pages 388-403.e18

(Bostrom 2003) Bostrom N. 2003. “Are You Living In a Computer Simulation?” Philosophical Quarterly. 53 (211): 243–255.

(Bublitz 2013) Bublitz J-C. “My Mind is Mine!? Cognitive Liberty as a Legal Concept.” In: Hildt E, Franke AG, eds. “Cognitive Enhancement. An Interdisciplinary Perspective.” Dordrecht: Springer; 2013. p. 233–64.

(Cannon et al. 2010) Cannon R.C., O'Donnell C. and Nolan M.F. 2010. ”Stochastic ion channel gating in dendritic neurons: morphology dependence and probabilistic synaptic activation of dendritic spikes.” PLoS Comput Biol 6: e1000886.

(Chalmers 2014) Chalmers D. 2014. “Uploading: A Philosophical Analysis.” In: Blackford R. and Broderick D. (eds). “Intelligence Unbound: The Future of Uploaded and Machine Minds,” p. 102–118. John Wiley & Sons, Inc

(Cole et al. 2008) Cole, B. S., Hopkins, C. M., Tisak, J., Steel, J. L., and Carr, B. I. 2008. “Assessing spiritual growth and spiritual decline following a diagnosis of cancer: Reliability and validity of the Spiritual Transformation Scale.” Psycho-Oncology, 17(2), 112–121

(Collins et al. 2025) Collins, L.-T., Huffman, T., Koene, R.A. 2025. “Comparative prospects of imaging methods for whole-brain mammalian connectomics.” Cell Reports Methods, 5(2)

(Citri and Malenka 2008) Citri, A. and Malenka, R.C. 2008. ”Synaptic plasticity: multiple forms, functions, and mechanisms.” *Neuropsychopharmacology*, 33(1), pp.18-41.

(Deary et al. 2021) Deary I.J., Cox S.R. and Hill W.D. 2022. ”Genetic variation, brain, and intelligence differences.” Molecular Psychiatry 27:335–353.

(Dongés et al. 2012) Dongés B., Haupt L.M., Lea R.A., Chan, Shum D.H.K. and Griffiths L.R. 2012. ”Role of the apolipoprotein E and catechol-O-methyltransferase genes in prospective and retrospective memory traits.” Gene, Volume 506, Issue 1: 135-140.

(Dorkenwald et al. 2023) Dorkenwald S., Matsliah A., Sterling A.R., Schlegel P., Yu S.C., McKellar C.E., Lin A., Costa M., Eichler K., Yin Y., Silversmith W., Schneider-Mizell C., Jordan C.S., Brittain D., Halageri A., Kuehner K., Ogedengbe O., Morey R., Gager J., Kruk K., Perlman E., Yang R., Deutsch D., Bland D., Sorek M., Lu R., Macrina T., Lee K., Bae J.A., Mu S., Nehoran B., Mitchell E., Popovych S., Wu J., Jia Z., Castro M., Kemnitz N., Ih D., Bates A.S., Eckstein N., Funke J., Collman F., Bock D.D., Jefferis G.S.X.E., Seung H.S., Murthy M. and FlyWire Consortium. 2023. ”Neuronal wiring diagram of an adult brain.” bioRxiv.

(Eliasmith and Trujillo 2024) C. Eliasmith and O. Trujillo. 2014. ”The use and abuse of large-scale brain models.” Current Opinion in Neurobiology, Volume 25, Pages 1-6.

(Fagella 2020) Faggella, D. 2020. ”Simulation #506 Dan Faggella - Programmatically Generated Everything.” Interview. https://simulationseries.com/ https://youtu.be/lyBGgE6wYR0, accessed 17 October 2022.

(Frankland et al. 2019) Paul W. Frankland, Sheena A. Josselyn and Stefan Köhler. 2019. ”The neurobiological foundation of memory retrieval.” Nature Neuroscience volume 22, pages 1576–1585.

(Fulton and Briggman 2021) Fulton, K. A. and Briggman, K. L. 2021. ”Permeabilization-free *en bloc* immunohistochemistry for correlative microscopy.” eLife 10, e63392.

(Goldstein and Hersen 2000) Goldstein G. and Hersen M.(editors). 2000. “Handbook of Psychological Assessment.” Elsevier. 3rd Edition, February 16, 2000

(Goldstein et al. 2003) Goldstein, G. (Editor), Beers S.R. (Editor) and Hersen M. (Editor-in-Chief). 2003. “Comprehensive Handbook of Psychological Assessment, Volume 1: Intellectual and Neuropsychological Assessment.” Wiley. July 2003

(Goriounova and Mansvelder 2019) Goriounova N.A., Mansvelder H.D. 2019. ”Genes, Cells and Brain Areas of Intelligence.” Front. Hum. Neurosci. 13:44.

(Gutzen et al. 2018) Gutzen R., von Papen M., Trensch G., Quaglio P., Grün S. and Denker M. 2018. ”Reproducible Neural Network Simulations: Statistical Methods for Model Validation on the Level of Network Activity Data.” Front. Neuroinform., 19.

(Hanson 2016) Hanson, R. 2016. ”The Age of Em.” Oxford: Oxford University Press.

(Hanuschkin et al. 2010) Hanuschkin A., Kunkel S., Helias M., Morrison A. and Diesmann M. 2010.”A General and Efficient Method for Incorporating Precise Spike Times in Globally Time-Driven Simulations.” Frontiers in Neuroinformatics. Vol 4:113.

(Herz et al. 2006) Herz A.V.M., Gollisch T., Machens C.K. and Jaeger D. 2006. ”Modeling single-neuron dynamics and computations: a balance of detail and abstraction.” Science 314(5796), pgs 80–85.

(Holt and Koch 1999) Holt G.R. and Koch C. 1999. "Electrical interactions via the extracellular potential near cell bodies." J Comput Neurosci . 1999 Mar-Apr; Vol. 6(2):169-184.

(Horsman et al. 2014) Horsman C., Stepney S., Wagner R.C. and Kendon V. 2014. "When does a physical system compute?" Proc. R. Soc. A 470: 20140182.

(IEEE 2019) IEEE Computer Society. 2019. "IEEE Standard for Floating-Point Arithmetic." 2019-07-22; pp. 1–84.

(Lappalainen et al. 2024) Lappalainen J.K., Tschopp F.D., Prakhya S., McGill M., Nern A., Shinomiya K., Takemura S., Gruntman E., Macke J.H. and Turaga S.C. 2024. "Connectome-constrained networks predict neural activity across the fly visual system." Nature 634, pgs 1132–1140.

(Linssen and Lemmens 2016) Linssen C. and Lemmens P. 2016. "Embodiment in Whole-Brain Emulation and its Implications for Death Anxiety", Journal of Evolution and Technology, Vol. 26 Issue 2, July 2016, pgs 1-15.

(Minerva 2021) Minerva F. 2021 "The Ethics of Brain Uploading", in Carissa Véliz (ed.), Oxford Handbook of Digital Ethics, Oxford Handbooks (2023; online edn, Oxford Academic, 10 Nov. 2021)

(Nader et al. 2000) Nader K., Schafe G. and Le Doux J. 2000. "Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval." Nature 406: 722–726.

(Koene 2012) Koene R.A. 2012. "Fundamentals of whole brain emulation: State, transition and update representations." International Journal of Machine Consciousness, 4(01), 5-21.

(Krzanowski and Krzanowski 2023) R. Krzanowski and J. Krzanowski. 2023. "Does whole brain emulation entail the emulation of mental disorders?" Argument Biannual Philosophical Journal 13(1):29-42.

(Müller and Indiveri 2015) L.K. Müller and G. Indiveri. 2015. "Rounding Methods for Neural Networks with Low Resolution Synaptic Weights." arXiv:1504.05767v1.

(Pearce 2023) Pearce, D. (wireheading.com). 2023. "Wirehead Hedonism versus paradise engineering." Retrieved Aug 8th, 2023.

(Petrovici et al. 2014) Mihai A. Petrovici, Bernhard Vogginger, Paul Müller, Oliver Breitwieser, Mikael Lundqvist, Lyle Muller, Matthias Ehrlich, Alain Destexhe, Anders Lansner, René Schüffny, Johannes Schemmel and Karlheinz Meier. 2014. "Characterization and Compensation of Network-Level Anomalies in Mixed-Signal Neuromorphic Modeling Platforms." PLoS ONE 9(10): e108590.

(Pfeil et al. 2012) T. Pfeil, T.C. Potjans, S. Schrader, W. Potjans, J. Schemmel, M. Diesmann and K. Meier. 2012. "Is a 4-bit synaptic weight resolution enough? – constraints on enabling spike-timing dependent plasticity in neuromorphic hardware." Front. Neurosci. Sec. Neuromorphic Engineering Volume 6.

(Rosenblueth and Wiener 1945) Arturo Rosenblueth and Norbert Wiener. 1945. "The Role of Models in Science." Philosophy of Science Vol. 12, No. 4 (Oct., 1945), pp. 316-321. The University of Chicago Press.

(Rothblatt 2013) M. Rothblatt. 2013. "Techno-immortalism 101: What are mindfiles?" In "Longevitize! Essays on the science, philosophy & politics of longevity", ed. F. Cortese, 62–66. Center for Transhumanity.

(Sandberg, 2014) Sandberg, A. 2014. "Ethics of brain emulations." Journal of Experimental & Theoretical Artificial Intelligence, 26(3), 439–457.

(Sandberg and Bostrom 2008) Anders Sandberg and Nick Bostrom. 2008. "Whole Brain Emulation: A Roadmap." Technical Report #2008-3

(Shiu et al. 2024) 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.-c., McKellar C.E., Sterling A., Costa M., Eichler K., Bates A.S., Eckstein N., Funke J., Jefferis G.S.X.E., Murthy M., Bidaye S.S., Hampel S., Seeds A.M. and Scott K. 2024. "A *Drosophila* computational brain model reveals sensorimotor processing." Nature 634, pp. 210–219.

(Sorokina et al. 2021) Oksana Sorokina, Colin Mclean, Mike D. R. Croning, Katharina F. Heil, Emilia Wysocka, Xin He, David Sterratt, Seth G. N. Grant, T. Ian Simpson and J. Douglas Armstrong. 2021. "A unified resource and configurable model of the synapse proteome and its role in disease." Nature Scientific Reports 11:9967

(Wang et al. 2019) Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. "GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding." Conference paper at ICLR 2019.

(Winding et al. 2023) Michael Winding, Benjamin D. Pedigo, Christopher L. Barnes, Heather G. Patsolic, Youngser Park, Tom Kazimiers, Akira Fushiki, Ingrid V. Andrade, Avinash Khandelwal, Javier Valdes-Aleman, Feng Li, Nadine Randel, Elizabeth Barsotti, Ana Correia, Richard D. Fetter, Volker Hartenstein, Carey E. Priebe, Joshua T. Vogelstein, Albert Cardona and Marta Zlatic. 2023. "The connectome of an insect brain." Science vol. 379, no. 6636

(Wybo et al. 2021) Willem AM Wybo, Jakob Jordan, Benjamin Ellenberger, Ulisses Marti Mengual, Thomas Nevian and Walter Senn. 2021. "Data-driven reduction of dendritic morphologies with preserved dendro-somatic responses." eLife 10:e60936.

(Xu et al. 2020) Xu, C.S., Pang, S., Hayworth, K.J., Hess, H.F. 2020. "Transforming FIB-SEM Systems for Large-Volume Connectomics and Cell Biology." In: Wacker, I., Hummel, E., Burgold, S., Schröder, R. (eds) "Volume Microscopy, Neuromethods". Vol. 155. Humana, New York, NY.