

The neuroconnectionist research programme

Doerig, A.*^{1,2}, Sommers R.*³, Seeliger, K.⁴, Richards, B.^{5,6,7,8,9}, Ismael, J.¹⁰, Lindsay, G.¹¹, Kording, K.¹², Konkle, T.¹³, van Gerven, M.A..J.¹, Kriegeskorte, N.¹⁴, & Kietzmann, T.C.^{1,2}

¹ Donders Institute for Brain, Cognition & Behaviour, Nijmegen, the Netherlands

² Institute of Cognitive Science, University of Osnabrück, Germany

³ Max Planck Institute for Psycholinguistics, Department of Neurobiology of Language, Nijmegen, the Netherlands

⁴ Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

⁵ Dept. of Neurology and Neurosurgery, McGill University, Montréal, QC, Canada

⁶ School of Computer Science, McGill University, Montréal, QC, Canada

⁷ Mila, Montréal, QC, Canada

⁸ Montréal Neurological Institute, Montréal, QC, Canada

⁹ Learning in Machines and Brains Program, CIFAR, Toronto, ON, Canada

¹⁰ Columbia University, New York, USA

¹¹ New York University, New York, USA

¹² Bioengineering, Neuroscience, University of Pennsylvania, Pennsylvania, PA, USA and CIFAR, LMB

¹³ Harvard University, Cambridge, MA, USA

¹⁴ Zuckerman Institute, Columbia University, New York, USA

* equal contributions

Corresponding author: adoerig@uos.de

Abstract:

Artificial Neural Networks (ANNs) inspired by biology are beginning to be widely used to model behavioral and neural data, an approach we call *neuroconnectionism*. ANNs have been lauded as the current best models of information processing in the brain, but also criticized for failing to account for basic cognitive functions. **We propose that arguing about the successes and failures of a restricted set of current ANNs is the wrong approach to assess the promise of neuroconnectionism.** Instead, we take inspiration from the philosophy of science, and in particular from Lakatos, who showed that the core of scientific research programmes is often not directly falsifiable, but should be assessed by its capacity to generate novel insights. Following this **view, we present neuroconnectionism as a cohesive large-scale research programme centered around ANNs as a computational language for expressing falsifiable theories about brain computation.** We describe the core of the programme, the underlying computational framework and its tools for testing specific neuroscientific hypotheses. Taking a longitudinal view, we review past and present neuroconnectionist projects and their responses to challenges, and argue that the research programme is highly progressive, generating new and otherwise unreachable insights into the workings of the brain.

Introduction:

While the study of cognition is a millennia-old endeavor (e.g., already present in Aristotle's *De Anima*), recent years have seen remarkable advances in both experimental and computational analysis techniques, yielding more powerful ways in which we can study and model computations in

the brain¹. Yet, the level at which cognition should best be understood remains a hotly debated topic. Modeling biology at the molecular level may not guarantee a deeper understanding of the core principles of cognition, any more than one brain can serve as an explanation of another. Instead, **the task of cognitive computational neuroscience is to find the right level of abstraction, close enough to biology to preserve the essential mechanisms, but abstract enough to discard details not required for cognitive function, while reproducing the trajectory from actively sensed input, through internal representations realized in neural processes, to complex goal-directed behaviors².**

Traditional experimental approaches, for example adopted in visual neuroscience, often operate at the rather coarse-grained explanatory level of contrasting experimental conditions. For instance, by running experiments with highly controlled stimuli, neural firing rates have been interpreted in terms of category selectivity: neurons are deemed selective for “faces”, “houses”, or “tools”³⁻⁶. This approach has merit. Its controlled settings allow for maximal interpretability and suggest a clear taxonomy of neural selectivity⁷. Yet, human-interpretable labels for neural activity are limited by the researchers’ imagination, or simply by labels that language makes available to us. Nature does not owe us easy answers: neural selectivity may often rely on more complex features that only imperfectly map onto human-interpretable categories⁸. In addition, **showing selectivity for a higher-level concept does not provide an answer as to how the brain may have arrived at this high representational level from noisy, lower-level sensory input data.**

Together, these observations highlight the need for neuro-computational models that can bridge explanatory and computational levels and predict neural data and behavior while being grounded in sensory data. This enables answering central questions of cognitive neuroscience: How can we link sensory input to neural data across brain regions, not only at the level of individual cells but also at the population level? How can neural processes be linked to behavior? How do neural representations change, not only through space, but also through time (from milliseconds to developmental trajectories)? How can past experience be encoded in the system and which types of feature selectivity allow for task-general robust performance?

In short, for a complete picture of how the mind works, neuroscience is in need of computational models that go beyond the limits of human-interpretable labels and tie together multiple levels of explanation. A good model of how some cognitive function is implemented in the brain should therefore (non-exhaustive list):

- a. Specify which computations are carried out by the brain (computational level).
- b. Show how these computations lead to complex behavioral patterns that can be tested in experiments (behavioral level).
- c. Show how these computations lead to complex neural dynamics that can be tested in experiments (single unit and collective dynamics levels).
- d. Show how these computations can be carried out in complex settings, beyond simplified experiments (rich domain knowledge).
- e. Show how these computations can be grounded in natural sensory information, rather than high level features provided by humans (sensory grounding).
- f. Show how these computations arise from adaptive processes that unfold at multiple timescales (inference, learning, development and natural selection).

To address these desiderata in a consistent explanatory framework, researchers have turned to computational modeling using Deep Neural Networks^{9,10} (DNNs; the latest incarnation of Artificial Neural Networks¹¹ (ANNs)). ANNs seem well-suited to tackle the challenges described above¹². First, ANNs are made of simple units that collectively implement complex computations that drive the network's behavior. That is, they offer a framework spanning the single unit, collective dynamics, behavioral and computational levels (desiderata a-c). Second, ANNs use millions (sometimes billions) of synaptic parameters to encode rich domain knowledge, while optimizing connectivity over time to learn (desideratum d). Third, **ANNs are grounded in sensory input, meaning that they can be trained on raw "sensory" data to fulfill "behavioral" needs, without the need for human-engineered input features, offering a link between sensation, cognition and action** (desideratum e). Finally, by allowing us to iteratively compare different biologically plausible learning rules and objectives, ANNs can help uncover how learning and cognitive development are made possible, and how they interact with architectural network features (desideratum f). Importantly, the architectural flexibility of ANNs and the different ways in which they can be trained allow for explorations, in a top-down fashion, of which biological details are needed to satisfy desiderata (a)-(f) for a given cognitive phenomenon. That is, **they allow for complex hypothesis testing, much like traditional experimentation**.

These characteristics allow researchers to rigorously test ANNs against large-scale behavioral and neural experimental datasets collected from a large array of brain regions¹³, and adjust the level of biological detail where needed - an approach markedly different from classic machine learning (ML) engineering, which is usually geared towards high performance on a small number of benchmarks. The approach furthermore goes beyond classic connectionist models of the 20th century¹⁴, which used smaller networks to explain higher-level cognitive tasks rather than the more recent focus on multi-level understanding of brain function. Due to the close integration into **neuroscience, both in terms of network design and mapping of internal representations to brain function and neural data, we term this approach *neuroconnectionism***.

Neuroconnectionism has already been fruitful in a wide variety of neuroscientific settings, including vision¹⁵⁻²¹, audition^{22,23}, semantics²⁴⁻²⁶, language^{27,28}, reading²⁹, decision making³⁰⁻³³, memory³⁴, game playing³⁵, motor control³⁶⁻³⁹, and the formation and coding principles of brain areas⁴⁰⁻⁴⁵. See section *The neuroconnectionist belt* for more details, and ⁴⁶⁻⁵² for reviews¹.

At the same time, neuroconnectionism does not remain unchallenged. It has been raised that ANNs differ strongly from biology, that they often behave in non-human ways⁵³⁻⁵⁶, and that the complexity of the models prohibits true understanding^{54,55}. These challenges are sometimes interpreted as suggesting that ANNs are not useful models for learning about the brain. For example, an argument might proceed like this: ANNs are sensitive to adversarial attacks, unlike humans; therefore, these models have been falsified and thus should not be used as models of the visual system. However, this logic only follows from a Popperian view of science (Box I).

Instead, we argue that these challenges are better seen as signposts. This changes the interpretation of the above example: **ANNs are sensitive to adversarial attacks, unlike humans; therefore, they must**

¹ Please note that the authors of this piece predominantly work in the sensory domain, in particular in vision. The provided examples are hence biased towards this research field. The larger neuroconnectionism community is however not limited to sensory processing, but involves, among many other areas, language processing, memory, (meta) learning, as well as movement and embodiment/robotics.

lack the relevant mechanisms that yield a more robust, human-like visual representational format, which will be important to improve future models. This conclusion follows a more Lakatosian philosophy (Box I). To elaborate on this view, we introduce neuroconnectionism as a Lakatosian research programme, which, with its explicitly longitudinal perspective on the developments in the field, offers fruitful grounds for discussion. This provides a new view on the challenges mentioned above, which act as signposts towards further developments.

Neuroconnectionism as a Lakatosian research programme

In philosophy of science, Lakatos⁵⁷ proposed a general framework to evaluate scientific approaches. According to his view, science is typically carried out within research programmes. Such programmes share a *hard core* of background assumptions that are not typically challenged from within the programme, and contain a *belt of* auxiliary hypotheses that are experimentally tested. While the core cannot be altered without abandoning the research programme, elements of the belt are and should be subject to change.

Given these two elements, core and belt, the value of a research programme is determined not just by its current experimental success relative to other research programmes, but also based on whether it is *progressive* rather than *degenerating* - an explicitly longitudinal perspective. Progressive research programmes generate new theoretical insights and novel predictions where some of these predictions are corroborated by empirical findings. Degenerating research programmes do not have these two characteristics: they often lack new theoretical developments and devolve into repeated corroboration of very similar ideas.

In the following, we will characterize and evaluate neuroconnectionism as a Lakatosian research programme composed of a set of auxiliary hypotheses in a belt centered around a unifying core, rather than a single falsifiable hypothesis. That is, **we consider productive ways to use ANNs to gain insight into understanding the mind and brain** (rather than considering any specific current model as *the* model of the brain that is true or false). We will describe the tools that are currently used for hypothesis testing in the field, review past challenges and successes, and revisit current controversies to evaluate whether neuroconnectionism is progressive.

Importantly, our arguments hold even if one does not accept the Lakatosian view on progress in science. While the Lakatosian distinction between core and belt is particularly well-suited to contribute to the ongoing debate about the merits and weaknesses of neuroconnectionism, what matters is that scientific theories always have core guiding principles that are relatively isolated from direct empirical testing - an uncontroversial view in the philosophy of science (see Box I). These core principles define directions of inquiry, including which experiments are conducted, and what is taken to corroborate or weaken theories. Theories are hence judged through a holistic assessment of their successes and fertility in guiding experimental pursuits. The Lakatosian perspective that we adopt is merely one helpful way of expressing these general ideas.

BOX I: THEORY SELECTION & PHILOSOPHY OF SCIENCE

Through much of the twentieth century, the Popperian view that theories are rejected when they are falsified in tests was dominant. A scientific theory generates predictions and tests are run to see if the predictions are correct. If they are not, the theory is falsified and can be rejected (it may take more than one test to show that the failure wasn't experimental error, but once the result is accepted, the theory must be rejected).

The Popperian view assumes that the logic of disconfirmation follows the following schema, where T is a theory, and O is an observation.

If T, then O

Not O

Hence, not T

It was noticed first by Pierre Duhem (the French theoretical physicist and historian of science) in 1906, and later reinforced by Quine (one of the most influential analytic philosophers in the twentieth century), that science does not work like that in practice and could not work like that in principle. One typically (and perhaps always) needs to combine the hypothesis to test with auxiliary beliefs to extract empirical predictions. When it is made fully explicit, the logic looks as follows:

If T, and A1, and A2, and A3, ..., and An, then O

Not O

Hence, not T, or not A1, or not A2, or not A3, ..., or not An

Where (A1, and A2, and A3, ..., and An) are auxiliary hypotheses needed to generate predictions.

For this reason, **it is never a single hypothesis, but a whole collection that generates predictions, any one of which might be at fault if the prediction is not vindicated.**

This fact about the logic of confirmation has been one of the centerpieces of twentieth century philosophy of science. Quine popularized the idea that a theory forms a web of beliefs related by inferential connections. The web has a topology with beliefs that can be most directly subjected to empirical tests at the periphery, and others that are insulated from direct testing by long chains of intermediary hypotheses. Beliefs at the periphery describe localized observable facts; beliefs at the center describe the kinds of general beliefs that guide the explanation of a whole body of phenomena. These are highly connected in the web and separated from empirical predictions by mediating propositions. One could hold onto the general beliefs in the face of mounting evidence if one was willing to make adjustments in other parts of the web.

Lakatos' distinction between core and belt acknowledges this holistic nature of confirmation, while offering a pragmatics of testing that is well-suited to capture how science actually works. In the Popperian view, testing is a matter of checking whether a theory accords with fact. On a Lakatosian view, it goes hand in hand with the development of theory. One adopts a set of core

principles as a kind of working hypothesis, proceeding on the assumption that they are correct and using them to try to understand the phenomena. Testing is a process that involves striving to bring theory and fact into closer agreement by exploring ways in which the core principles can be preserved while accommodating the evidence. If the core principles are held fixed and leeway is granted to explore alternative auxiliary hypotheses, testing can be directed at the belt, giving the theory every chance of preserving the core while accommodating the evidence. This is a way of targeting disconfirmatory evidence against one part of the web. A theory is rejected not as the result of a direct conflict with the evidence, but because the attempt to preserve the core principles becomes so cumbersome that they cease to form a productive working hypothesis.

For example, in astronomy, there were deviations in planetary trajectories from the smooth ellipses predicted by Newtonian mechanics. Instead of rejecting Newtonian laws, people interpreted this as a reason to look for ways to explain the distortion, assuming the correctness of the laws (e.g., an unseen planet). Hence, a belt claim was falsified but the core was not abandoned. Instead, the core is changed only when it becomes unproductive to hold onto it. For example, in the twentieth century, evidence accumulated against Newtonian celestial mechanics that could not be solved assuming the correctness of the laws, which led to its rejection and the development of general relativity.

The neuroconnectionist core

As discussed in the introduction, ANNs hold promise to fulfill desiderata (a)-(f) by tying together multiple levels of explanation, incorporating rich domain knowledge, and going beyond the limits of human-interpretable labels. In a sense, neuroconnectionism can be seen as an attempt to address these desiderata by using ANNs, which motivates why ANNs are at the core of the research programme. In this section, we elaborate on this further and describe the core tenets of neuroconnectionism and how they relate to specific properties of ANNs.

ANNs consist of interconnected sets of rather simple units, which perform local computations and which by themselves have no perceptual or cognitive abilities. Instead, perceptual or cognitive computations emerge as a collective property of the network, distributed across many units. Information flow through the network determines its function, and is itself determined by the connection strength between the many units. How to set up and weigh these connections to encode domain knowledge and perform according to a set objective is a difficult problem: **for a fully connected recurrent ANN with N units, there are $2^{N \times N}$ possible structural connectivity patterns, under the simplifying assumption of connections being either 0 or 1².**

Because of this complexity, researchers constrain the connectivity motifs in the form of network *architectures* that determine which connections are possible - providing strong inductive biases. Most current network architectures organize their units in layers and within-layer-connectivity is determined as part of the architecture. In feedforward networks, units of a given layer only receive inputs from the layer below and only send outputs to the layers further above. Recurrent networks

² For a system of 8 units, that is 18 quintillion 446 quadrillion 744 trillion 73 billion 709 million 551 thousand 616 connectivity patterns.

allow for additional top-down and lateral message passing. Another noteworthy architectural aspect to constrain the number of learnable parameters in a network is to make it convolutional^{16,58,59}. This convolutional architecture was directly inspired by seminal findings in visual neuroscience⁶⁰. Convolutional layers are subdivided into feature maps. Each unit in a feature map has its own receptive field, but all units share the same selectivity. Applied in the visual domain, this means that when the network learns to detect a feature at one location, it automatically generalizes to other locations. This is a useful inductive bias for vision: since the same object can occur anywhere in an image, it is sensible to look for the same features everywhere, without having to re-learn them at each location. Other domains, such as olfaction, do not share this property and thus favor other inductive biases. Many different architecture types exist in the literature, but the above examples illustrate the main point: by constraining connectivity, each architecture not only reduces the number of parameters in the network, but also implements inductive biases that determines which functions can be efficiently learnt.

While architectural design defines the general space of possible solutions, the individual network parameters still need to be optimized. For example, ResNet-50, a popular convolutional architecture, has over 23 million parameters⁶¹. To learn these parameters, networks are typically presented with input *datasets* (e.g. natural images, auditory signals, or text corpora) and their target behavior is defined based on one or several mathematical *objectives* (also called loss functions). Ultimately, several low-level objectives may be subsumed in higher level objectives, such as “fitness”. Deviations from the target, or errors, are then minimized by updating the weights between network units, which is determined by a *learning rule*. To do so, the weights of each individual unit must be changed in accordance with the contribution of the unit to the whole network’s error. Given that ANNs can have millions of units (operating over extended time in the case of recurrent networks), finding good learning rules is far from trivial. How to attribute the contribution of each unit to the overall error is called the *credit assignment problem*. The most widely used learning rule to solve credit assignment in ANNs is backpropagation⁶², but there are many others too, e.g. ^{63,64}. It is worth noting that each learning rule comes with its own hyperparameters, such as a learning rate and how the learning rate changes over training. Overall, many datasets, objectives and learning rules are used in neuroconnectionism, and they strongly impact the behavior and internal states of ANNs (see section *The neuroconnectionist belt* for details).

The approaches described above for designing and training ANNs offer ways of testing hypotheses of how certain functions may be implemented in the brain while enabling researchers to perform large-scale in-silico experiments^{46,49–51}. Indeed, an ANN’s architecture, dataset statistics, objectives, and learning rules can be mapped onto central questions of cognitive neuroscience, allowing computational traction to untangle the interacting contributions of pre-specified structure (for ANNs: determined by the architecture) and experienced input (for ANNs: training dataset), why neural selectivity in a given brain region is the way it is (for ANNs: which objectives are being optimized), and how the brain may adjust its internal representations (for ANNs: credit assignment / learning rules). All of these questions can be studied across levels of explanation and temporal scales, incorporating rich domain knowledge grounded in sensory data, in line with desiderata (a)-(f).

Importantly, there is likely no shortcut to the iterative learning and processing occurring in neural networks. Many complex iterative processes – such as cellular automata like Conway’s Game of Life or Rule 110 – can be specified on half a post-it note, yet their eventual behavior after many steps is

fundamentally unpredictable^{65,66}. Similarly, the learning and processing dynamics of neural networks are unpredictable despite having simple mechanics. This unpredictability has been proven for recurrent neural networks⁶⁷. Put differently, **ANNs need to be processed iteratively to study the effects of training, architecture, task, and dataset. No analytical mathematical model exists to date that could use shortcuts to directly get the final result without going through the iterative process.** The same is likely true of the brain, a complex, deep, massively parallel and recurrent neural network. Therefore, it likely can only be modeled using another complex iterative process, such as ANNs, while performing careful hypothesis testing at a level of abstraction that offers the best insights to the process at hand (the optimal level of abstraction may depend on the phenomenon under investigation). This lack of simpler models justifies the use of complex ANN models in neuroconnectionism.

With this information in place, we can state a Lakatosian core of the neuroconnectionism research programme (see also Figure 1):

The core of neuroconnectionism is that brain computations, representations, learning mechanisms, and inductive biases are best understood via modeling in ANNs that are defined by their architectures, input statistics, objective functions and learning rules. Multilevel and sensory-grounded iterative models like ANNs are needed to model multilevel and sensory-grounded iterative brain computations, and thereby achieve desiderata (a)-(f).




Figure 1 | Illustration of the neuroconnectionist research programme. The core of neuroconnectionism (depicted as the gray center) is to use ANNs, defined by their architectures, objective functions, input data statistics and learning rules, to test neuroscientific hypotheses. Following the Lakatosian view, experimental investigations are not directed at the core of the research programme, but at a surrounding belt of auxiliary hypotheses (depicted as coloured arcs, e.g., the green arc could stand for the hypothesis of CNNs as models of the human vision). Falsifying individual belt hypotheses does not invalidate the research programme as a whole, but instead calls for improving the belt. Hence research programmes are not judged by attempting to falsify the core directly, but rather by assessing if they are progressive or degenerating. Good research programmes are progressive: they successfully integrate hypotheses into a coherent belt, produce new hypotheses, and make testable predictions that are corroborated.

The above definition aims to capture the core ideas shared by the many projects that contribute to neuroconnectionism. However, like all research approaches, there can be no rigid definition. Still, just

like how members of a family do not share *all* traits, but still possess a family resemblance⁶⁸, so too do members of the neuroconnectionism research programme constitute a diverse but cohesive family of approaches. Distinctive features of neuroconnectionist models include:

Explaining cognition as primary explanatory goal: Neuroconnectionist models are primarily aimed at explaining cognitive functions, not at describing biology with the highest possible detail. Biological detail is added to models in a top-down, hypothesis driven fashion, when these details are necessary for explaining behavioral or neural data. This makes neuroconnectionism different from approaches aiming to perfectly replicate a human brain in silico and from biophysical models aiming to model every aspect of a neuron.

Distributed representations and computations: The modeled property emerges from the collective behavior and dynamics of simple units which, taken independently, do not show the modeled property. This distributed nature is central to neuroconnectionism, since it naturally bridges between explanatory levels, from single units, through collective dynamics and onto behavior. This distinguishes neuroconnectionism from traditional models where each parameter has a straightforward interpretation, such as drift-diffusion models, and requires special interpretation frameworks to cope with the distributed nature of computations (see the *neuroconnectionist toolbox* section).

Iterative training & inference: The high-dimensionality of distributed models makes it virtually impossible to tune all parameters by hand, so iterative training is required. At inference too, the behavior and dynamics cannot be simplified into a simple interpretable equation: to know what the model predicts, one needs to run the model. This iterative nature distinguishes neuroconnectionism from approaches that can be formulated in a small number of non-iterative equations, such as (hierarchical) Bayesian models.

There are of course many edge cases. For example, while grounding in sensory input is a desideratum of the neuroconnectist approach, we note that not all neuroconnectionist models are grounded in sensory input. In language or memory models, for instance, the input to the model often consists of high-level concepts, such as words. Still, we consider them neuroconnectionist, since the sensory nature of the inputs may not be relevant in these particular cases, and could be added to later models if needed. As another example, certain neuroconnectionist models may not explicitly employ a “behavioral” objective function, but rather attempt to minimize energy efficiency (e.g. ⁶⁹). End-to-end learning where models are directly fitted to neural data instead of trained on a behavioral task are also neuroconnectionist since they are a way to test hypotheses about which architectures are capable of reproducing neural dynamics¹⁹. In addition, models do not need to model cognition in general, but can focus on specific components. Understanding which aspects of cognition can only be modeled jointly is part of the hypothesis testing process. For example, sensory-motor interactions in embodied models may or may not be required to explain certain aspects of visual processing. Lastly, models that are architecturally different from the brain, such as GANs or Transformers, can also be used to test neuroscientific hypotheses and thereby contribute to neuroconnectionism⁷⁰⁻⁷². In short, **most ANN models that can be used for neuroscientific hypothesis testing can be considered neuroconnectionist given a researcher’s motivation of how the model could be rooted in biology, and how it tests theories about neural computations.**

Importantly, neuroconnectionist approaches are about delineating mechanisms and theoretical understanding through hypothesis-driven research, which differs from the goals of engineering. Around the time of the first large-scale vision networks^{58,73}, neuroconnectionist models were borrowed directly from engineering applications. For example, it was found that the engineering models that performed the best on engineering benchmarks (as of 2014) were also better at predicting brain activity²¹. This may have led to a form of “computational opportunism” in which researchers could directly test ML models against neuroscientific data without having well-formed hypotheses. However, this trend has come to a halt with more recent network architectures where better task performance does not align with improved predictions of neural firing rates¹³. Indeed, when taking the hierarchical nature of the visual system into account, recent years have seen a constant *decrease* in agreement between ANNs and neuroscience⁷⁴. These trends highlight that **focusing on engineering goals based on performance alone is not sufficient for a neuroscientific understanding. Neuroconnectionism needs to develop its own models and metrics.**

Following Lakatos’ definition of research programmes, the different neuroconnectionist projects all use ANN models as core to their approach. Experiments are not directly applied to the core, but rather to the surrounding belt hypotheses, which are realized via ANN instantiations and tested using various techniques. The promise of the overall research programme is then to be judged based on its ability to generate new insights, corroborate belt hypotheses, derive new testable predictions, and, perhaps most importantly, productively address criticisms to make progress.

The neuroconnectionist toolbox

To implement and evaluate the neuroconnectionism research programme, tools are needed to instantiate and test models in its belt. These include tools to test to which extent neural network model instances align with neural data and behavior (Figure 2). Empirical findings revealing areas where ANNs fall short are an essential part of the framework and approach, as they provide crucial data points indicating where our understanding and models need to be improved.

Network design and training

A central aspect of the research programme is the creation and training of ANNs. For this aspect, engineering has been instrumental, as it provided researchers with key hardware and software technologies that enable training of task-performing models.

To design ANNs, researchers can choose from a large zoo of network architectures and unit types. The different unit types are abstracted away from the brain to various extents. For example, common computational units go from very simple rectified linear summation (ReLU), to complex units modeling basic memory (e.g. LSTMs). Importantly, researchers have started integrating various aspects of biology into network design while testing them against neural and behavioral data^{13,75}. This includes recurrent connectivity^{19,76–82}, and different base-units, such as richer rate-based neurons^{83–85}, spiking neurons^{86–89}, and neurons with multiple compartments^{80,90,91}.

To train ANNs, researchers need to choose datasets, objectives and learning rules. Many large-scale datasets are openly available. Typically, external datasets are used (i.e., the dataset is the input to the network and the network's output is trained to match the labels). For example, inputs could be images and outputs could be captions giving a semantic description of the image. In addition, networks can also be directly trained to mirror brain activity^{19,77,92-99}, leveraging recent efforts to record large neural datasets¹⁰⁰⁻¹⁰⁵. The dataset determines the input statistics that the network learns from, so an important avenue of research is to develop more naturalistic datasets (e.g. ¹⁰⁶).

A large variety of objectives is used, including supervised (e.g. classification, scene captioning), unsupervised (e.g. contrastive learning^{107,108}, predictive coding^{109,110}, image generation^{44,111,112}, temporal stability¹¹³⁻¹¹⁷, and energy efficiency⁶⁹) and behavioral reward^{30,35,118}. The objective determines what the networks try to learn based on the input statistics. Different objectives may for example lead to modeling different brain areas¹¹⁹. As another example, one may test whether the ventral visual stream is better modeled using an object categorization or a semantic objective²⁴, or test the impact of noise in the dataset¹²⁰.

The most common learning rule is, by far, backpropagation⁶². In its standard form, backpropagation is not biologically plausible, and a lot of work seeks more plausible versions^{91,121-124}. Other learning rules exist too^{63,64,125}. Each algorithm has its own characteristics. For example, gradient descent tends to learn the input features with most variance first¹²⁶ and produces efficient neural codes¹²⁷. Hebbian learning, on the other hand, is a simpler local rule that has been directly observed during learning in biological systems¹²⁸.

In summary, by using frameworks optimized for big data and high throughput, computational neuroscientists can create ANNs and test which biological aspects are necessary to reproduce brain function and behavior. ANNs thereby offer a top-down approach to understanding which aspects of the brain are central to its function.

Model testing

Although critics sometimes claim that we cannot understand ANNs^{54,129}, **ANNs are not black boxes. Rather, they are transparent boxes with easy access to each unit's activities and connections.** Researchers can perform in-silico experiments at speeds that are orders of magnitude faster than on any biological brain, and, for now, free of ethical concerns that come with classical experimentation. Several tools have been developed to improve the interpretability of ANNs and test them against biological data.

Hypothesis testing in neuroconnectionism relies on testing trained ANNs against brain data on various levels, from neural data up to behavioral patterns. Because of multiple realizability, systems that are structurally very different from the brain can nevertheless enable successful predictions of human behavior, or of the activity in certain brain areas (e.g. ^{130,131}). Hence, to ensure that an ANN implements a given cognitive function in a human-like way, it needs to map onto human brain processing across levels, ranging from behavioral to neural data, ideally while taking into account the physical constraints that the brain faces, such as, for example, metabolic costs^{69,132-134} and considerations of wiring length⁴⁵. Importantly, no single method is perfect, and various

complementary approaches are needed. Hence, **developing good metrics to compare humans and ANNs across levels is an essential part of neuroconnectionism.**

Behavioral agreement: At the behavioral level, the outputs of ANNs can be compared to human responses in several settings. Coarse measures such as overall task performance are useful but only of limited precision, since ANNs are becoming increasingly good at tasks for which humans were until recently deemed the gold standard. These include object recognition^{58,61,135}, board games¹³⁶ or video games^{137,138}. To overcome these limitations, several more fine-grained methods exist. These include the use of diagnostic readouts to characterize the information represented in a population of units and to translate it into behaviorally relevant measures such as reaction times¹³⁹, detailed analysis of error patterns¹⁴⁰, testing on out of distribution examples¹⁴¹⁻¹⁴³, and reproducing psychophysical results that target specific aspects of processing^{76,144-150}. See ¹⁵¹⁻¹⁵⁴ for further discussion about how to compare human and ANN behavior. A unified model addressing years of psychophysical experimentation is an important target, yet to be achieved.

Agreement with neural data: At the neural level, the activity patterns of ANNs can be compared to biology in several ways. One common way of doing so is using **Representational Similarity Analysis^{155,156} (RSA)**. **RSA characterizes the internal representations of a system by quantifying the dissimilarities between the population responses to various stimuli.** Internal representations of ANNs and the brain are deemed similar if the corresponding response geometries agree. This approach has the benefit of side-stepping the need to directly map from individual ANN units to individual neurons or voxels. It is therefore well suited for comparing ANN activations with neural data, e.g. obtained via neuroimaging or electrophysiology. RSA can be applied using all units of a whole network, or all units of a network layer, units in a feature map, or individual ANN units. The representational geometries measured by RSA can be compared directly between different brain regions and/or ANN layers¹⁰⁶, or they can be combined using linear reweighting to optimally map onto each other^{18,157,158}. Interestingly, network units can be multiplied with a constant factor while not affecting network output if the synaptic connections to downstream units are downscaled accordingly. This implies that networks with identical behavior populate a whole subspace of solutions, which still differ in their alignment with brain data if no reweighting is applied. **An important challenge for current RSA methods is that they sometimes fail to discriminate between different network architectures, which all perform similarly on this metric¹⁵⁸.** Recent work has started addressing this issue by improving current RSA methods¹⁵⁹ and clarifying which aspects of brain computation should be targeted by RSA¹⁶⁰.

In addition to RSA, which is predominantly aimed at characterizing responses at the population level, encoding models can be used to predict the activity of single neurons or voxels across a range of conditions^{21,161,162}. Here, the activity of each biological unit (neuron or voxel) is predicted as a linear combination of ANN unit activations. To prevent overfitting, the underlying GLMs are typically regularized. More recent encoding methods separate spatial from channel dimension activity^{96,98}. In their original form, encoding models are not constrained as to which units can explain which biological counterpart. In principle this implies that, for example, higher-level brain regions can be explained by lower-level network features, or that thousands of cells in the brain can be explained by the response of a single network unit with broadly similar selectivity. New developments are underway to include an **ordered hierarchical mapping from ANNs to the brain^{74,163}, providing a more rigorous test of the alignment between brain representations and ANN models.**

A third approach, which builds on encoding models, is to use ANNs to design stimuli that strongly activate a single target biological neuron or area¹⁶⁴⁻¹⁶⁶, allowing for a causal interaction between ANNs and biological brains. This approach therefore goes beyond the otherwise correlational RSA and encoding model approaches.

In-silico electrophysiology: In addition to estimating the level of agreement between ANNs and biological brains in terms of behavior and neural recordings, ANNs themselves can be experimented on to better understand their inner workings. Since we have immediate access to all units, their activities and connectivity, almost any experiment is possible. This includes reliance on network initialisation¹⁶⁷, representational similarities across network architectures¹⁵⁹, and tests for the emergence of brain-like computations⁶⁹, selectivity profiles^{41,45} and cell types¹⁶⁸. To this end, standard neuroscientific methods can be applied, such as searchlight decoding, measures from signal detection theory, tuning curve analysis, and many more. In addition, ANNs can be selectively lesioned to test the impact of different parts of the network on its ability to map onto brain function. For example, the effect of recurrent connections can directly be assessed by ablating them^{19,82,147}. In-silico lesion studies are of course not limited to analyses of networks on their own, but can also be used to evaluate the changes in agreement between network and neural or behavioral data. Lastly, the ability to replicate topographic elements of brain organization can be tested in ANNs to better understand the origins and functional implications of such representational arrangements. First forays in this direction are well underway^{41-43,45,169,169,170}.

Developmental agreement: All of the above methods can be applied at different points during network training, from untrained to fully trained models, and these learning trajectories can be compared with different stages in biological development¹⁷¹. Although it is currently unclear which aspects of learning in ANNs are better seen as corresponding to learning during evolution and which are better seen as modeling learning during an organism's lifetime¹⁷², both can be addressed experimentally. **For example, the order in which children learn different words can be predicted by the performance of ANNs trained on visual classification and captioning tasks, over and above the expected effects of word frequency**¹⁷³.

In summary, neuroconnectionism has at its disposition a large array of techniques to train and evaluate neural networks as models of brain computation. Evaluating and contrasting different models for their ability to explain brain data are vital aspects of the research programme as they enable rigorous hypothesis testing, which brings developments in the belt of hypotheses, as discussed next. The fact that each model can be extensively tested from single neurons to behavior is a strong asset of neuroconnectionism that arguably no other approach can claim.

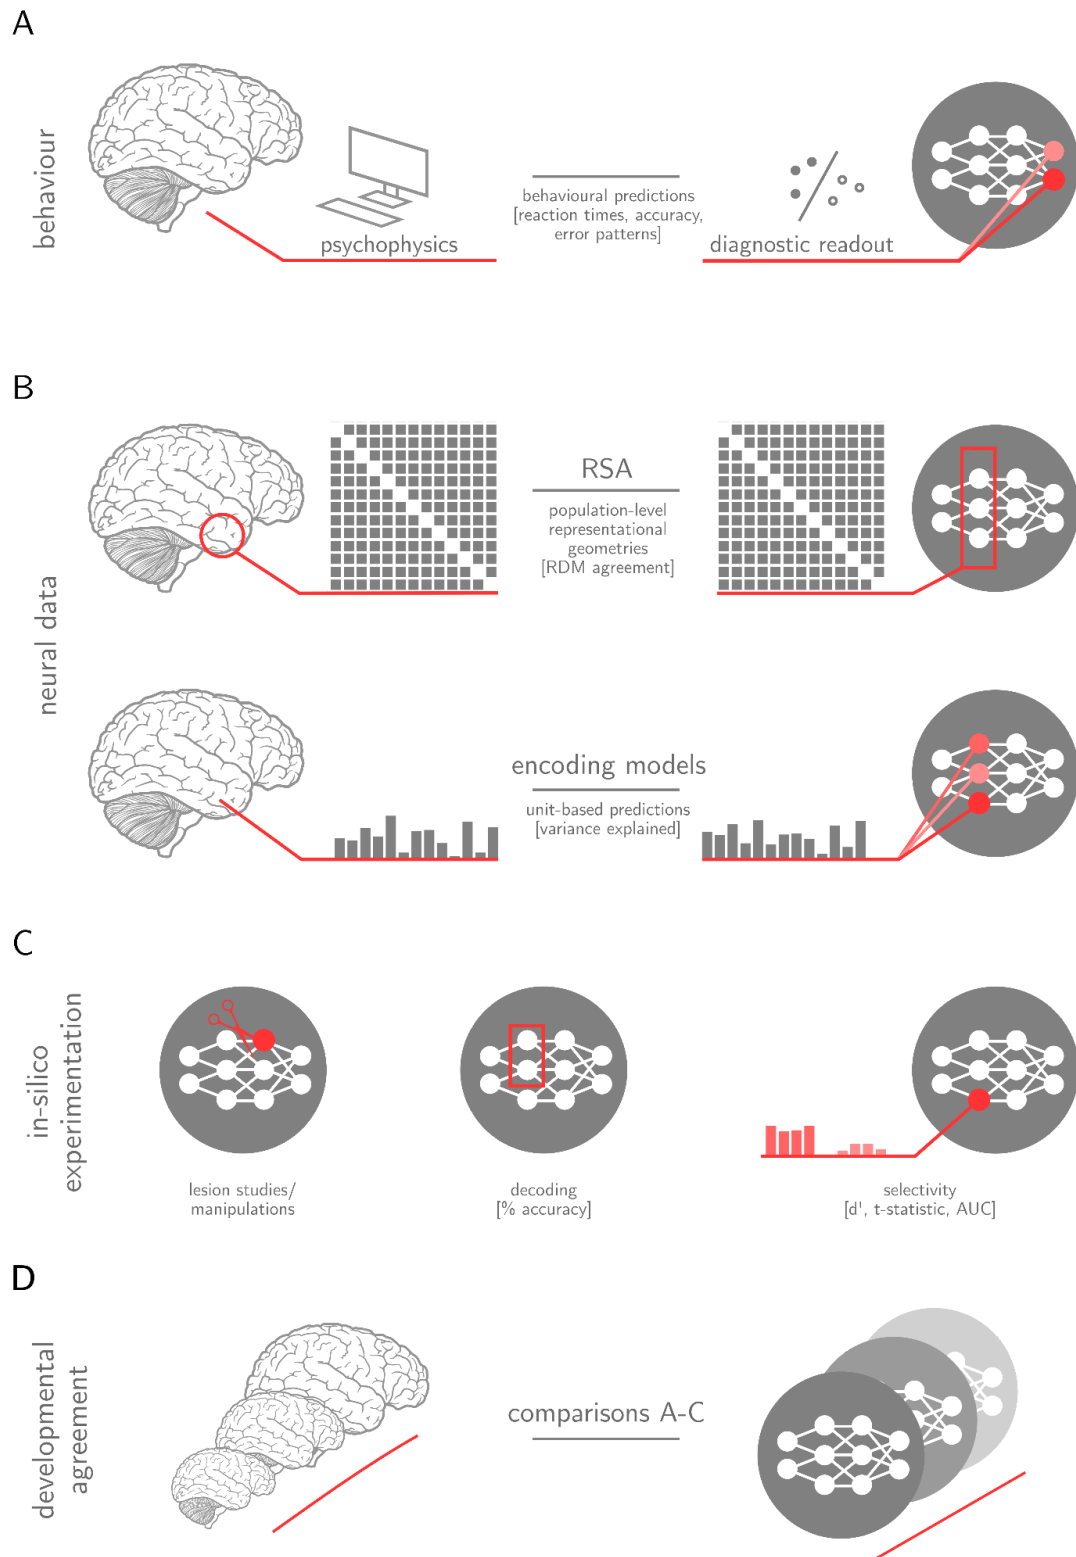


Figure 2 | The current neuroconnectionist toolkit for model testing. Similar to the belt of a research programme, its toolkit is under constant development. Next to improvements of existing techniques (which historically have focused heavily on sensory processing), this involves new approaches to widen the scope of phenomena that can be investigated. These methods can be applied to compare ANNs with any species (human, primate, mouse, ...). **A** Behavioral agreement. Outputs of ANNs are compared to human responses in

diverse settings, such as classification errors and accuracy, reaction times, action patterns, and others. **B** Agreement with neural data. Presenting identical stimuli (input) to the brain and computational model, the recorded brain activity patterns are directly compared to ANN layer activity patterns. The most common methods are RSA (comparing representational geometries) and encoding models (pointwise linear regression). **C** In-silico electrophysiology. ANNs are studied as in-silico models of cognitive functions with standard neuroscientific methods, such as manipulations and lesions, information decoding, unit-based tuning functions and others. **Effects of design choices such as recurrent connections, or manipulations such as detailed lesioning patterns can be studied extensively in this setup, something not possible in-vivo.** **D** Developmental agreement. Comparing the above at different stages of training to stages of learning in biological brains permits insights into cognitive development. Examples include behavioral patterns, map formation, or changes in neural selectivity with visual experience.

The neuroconnectionist belt

We have now defined the core of the research programme, as well as tools that are used to design, train, and evaluate ANN models across various levels of explanation. This brings us to the belt of the research programme, a set of hypotheses which are to be tested and which evolve as new evidence is integrated. Individual elements of the belt are important, but a more central aim, when taking a Lakatosian perspective, is an evaluation of longitudinal developments (theoretical and empirical), which determine the state of a research programme as either progressive or degenerative. By definition, the belt of a research programme is subject to change. New hypotheses can be derived, and existing hypotheses can be corroborated, altered, and rejected. An individual belt hypothesis that is rejected does not refute the core assumptions upon which a programme is built, but rather provides an important datapoint for future developments. According to this view, the overarching question becomes: How does the neuroconnectionism research programme fare in terms of productivity? Does it generate new insights, and does it address existing challenges? Are there challenges that cannot be overcome by the research programme (roadblocks: rendering it degenerating), or are challenges rather signposts towards open questions that can be addressed by improving the research programme (progressive)?

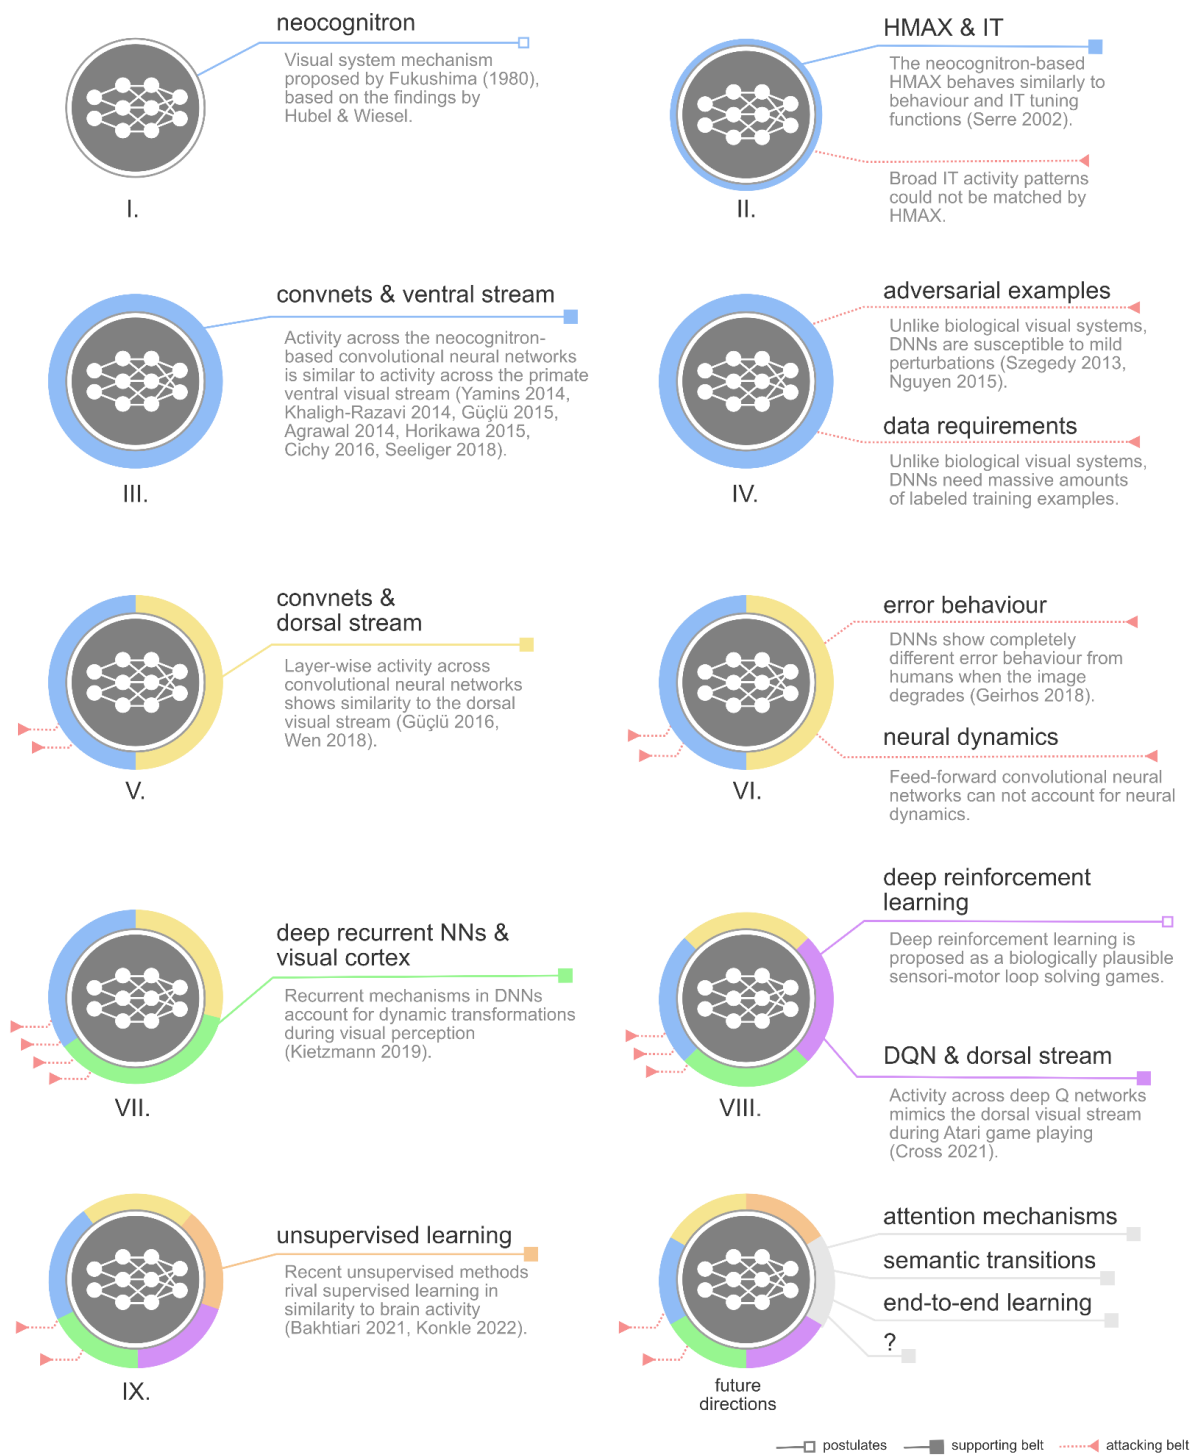


Figure 3 | Summarized historical progression of the neuroconnectionism belt in the area of visual computational neuroscience. The figure shows how empirical and analytical findings are integrated into (or strengthen) the belt, and how objections are processed within a constructive, productive research programme. **I.** The neocognitron is suggested as a mechanism behind the seminal findings about simple and complex visual system cells by Hubel & Wiesel. It can learn and recognize increasingly abstract visual patterns through mechanisms that are similar to convolutions. **II.** HMAX, a more powerful model similar to the neocognitron, is investigated closely for its similarity to neural activity. It matches IT tuning functions but does not match its broad activity patterns. It can not recognize objects in naturalistic photos. The neuroconnectionist belt

dwindles. **III.** Convolutional neural networks are successfully trained on large collections of naturalistic photos. A series of studies shows that their layer activity matches neural activity patterns along the primate ventral stream. This is the first time an image-computable and functional object recognition network matches activity patterns across the ventral visual system. These findings strengthen the neuroconnectionist belt and spawn a massive series of new studies and postulates. **IV.** With their susceptibility to adversarial attacks and the amounts of labeled training data they require, convolutional neural networks are shown to exhibit several behaviors that are not biologically plausible. **V.** Convolutional neural networks show similar layer activity to the dorsal stream, adding further experimental evidence strengthening the belt. **VI.** Adding to the growing list of biologically implausible mechanisms attacking the belt, it is shown that the error behavior during image degradation clearly diverges between humans and convolutional neural networks. Furthermore, feed-forward convolutional neural networks embody too simple mechanisms to cover neural dynamic observations beyond coarse rate coding. **VII.** The neural dynamics objection is resolved by the demonstration that dynamic transformations during visual processing can be captured if recurrence is added to DNNs. As recurrence is a biological property of the brain, this strengthens the neuroconnectionist belt. **VIII.** It is shown that activity across the dorsal stream during game playing matches activity in deep reinforcement learning networks, which implement a sensori-motor loop for the same game playing tasks. **IX.** Large numbers of labeled data are not required by unsupervised learning. Newer developments in this field are shown to match brain activity similarly well to supervised learning. As unsupervised learning is considered biologically plausible the resolution of this objection further strengthens the belt. **Future directions.** Attention mechanisms, semantic objectives and end-to-end learning in which networks are trained directly to match neural activity are all recent developments in ANNs. Experiments will tell which brain processes are better modeled by incorporating these elements.

The evolution of the neuroconnectionism belt: a progressive generator of new insights.

The belt of neuroconnectionism has considerably evolved in the past decade. By rapidly testing and expanding belt hypotheses, the community now commonly explores different architectures, datasets, objectives, and learning techniques across various experimental settings to test the resulting models for their alignment with brain and behavioral data.

One of the clearest examples of this progressive evolution in the field has been the change in how vision is modeled with ANNs in recent years (see Fig. 3 for an illustration of the historical progress of neuroconnectionism in vision). In the early 2010's, researchers focused a great deal of effort on comparing neural and behavioral data to what were, at that time, the state-of-the-art in ANNs for vision, namely, deep convolutional neural networks trained on image classification tasks^{15-18,20,21}. Over time, as researchers explored the successes and the failures of these models, the field has seen almost every component of these initial models updated, leading to new models that do a much better job of accounting for neurophysiological and psychological data. The earliest such update was a result of the recognition that the feed-forward nature of standard convolutional networks was both obviously incongruent with real neural anatomy, and functionally limiting. It was shown that recurrent networks were better at matching both behavioral data and neural activity patterns that occur with longer delays^{19,33,76-80,82,150,174-177}. Additionally, researchers explored ways to improve the training datasets beyond computer vision benchmarks, including a more ecologically relevant selection of object categories¹⁰⁶, video data¹⁷⁸, embodiment^{179,180} and goal-directed eye-movements¹⁸¹. Similarly, the use of supervised category training, which was always problematic from a neuroscience perspective - since humans do not need millions of labeled examples to learn - was shown to be unnecessary: recent work has shown that self-supervised techniques for training

ANNs lead to as good or better matches to neural representation and animal behavior^{108,182,183}. Moreover, self-supervised training on video data can account for the distinction between the dorsal and ventral pathways in the brain¹⁸⁴ and self-supervised training on lower resolution inputs provides a better fit to the mouse visual cortex¹⁸⁵. Other losses have also been explored, and researchers have found that dorsal stream vision can be explained both by control-based optimisation³⁵ and by self-motion related losses¹⁸⁶. A similar evolution has occurred in models of the hippocampal formation and related networks, where initial architectures and losses have been replaced as the belt of the research program evolved. Early connectionist models largely focussed on attractor networks^{187,188}, which captured many interesting aspects of these circuits. But, with time, these models have evolved to incorporate additional architectural features and loss functions related to prediction and spatial integration, leading to better and better matches with a host of experimental results^{189–191}.

At the same time, as new ANN architectures have come out with more sophisticated attention mechanisms, such as transformers, more and more research has demonstrated that transformers trained in a self-supervised manner can effectively capture the representations observed in language areas of the brain^{27,28}, and other circuits, such as the mnemonic circuits of the medial temporal lobes. This leaves open the possibility for numerous other updates and explorations to existing models of vision and other senses, including the use of self-attention, and the use of multi-modal networks that combine linguistic inputs with vision or other sensory modalities^{192,193}.

In the above examples, neuroconnectionist models have provided state of the art models of behavioral and neural data, discovering new perspectives on the roles of known processes such as recurrence, attention and the dorsal stream. Importantly, fundamentally new theoretical insights and predictions about brain processing have also been derived, and some have been directly corroborated by in vivo experiments (see Box II).

BOX II: PREDICTIONS & INSIGHTS GENERATED BY ANNs

An important aspect of neuroconnectionism, as with all of computational modeling, is that it is not limited to providing accounts of existing data but can propose, test, and discover fundamentally new principles. In the following we describe some of such findings.

ANNs have been used to derive novel predictions about neural selectivity. Using optimization techniques that work at the level of the network input, it was shown that ANNs can be used to design novel stimuli that drive biological neurons in a highly specific manner, exceeding the firing rates observed using natural stimuli^{164–166,194,195}. This causal manipulation provides an important extension to otherwise correlational tests of the alignment between models and the brain.

Another example of neuroconnectionism as a key for novel, and testable predictions about neural selectivity was provided by Bao et al.⁴⁰ who used ANNs to derive a theory of object selectivity in macaque IT. While certain parts of macaque IT had well-known selectivity profiles (e.g. face patches⁶), other parts remained uncharacterised “no-man’s land”, without any clear selectivity pattern. By using dimensionality reduction techniques on the late layers of an ANN trained on object classification, Bao et al. obtained a low-dimensional “object space” and found that responses of IT cells to a large set of objects aligned well with the axes of this space, including the

previously uncharacterised no-man's land units. With this, the authors were able to provide a unified picture of IT organization as following a low-dimensional object space, as extracted from a deep network. See ^{31,44} for further such discoveries.

Neuroconnectionism has also yielded various new insights into the computational basis of neural information processing. For example, based on the learning dynamics of ANNs, Saxe et al.²⁶ derived a mathematically tractable theory of semantic learning that recapitulates many empirical phenomena in human semantic development. Averbeck¹⁹⁶ proposed to explain adolescent changes in working memory by pruning in ANNs, and Rust & Jannuzi¹⁹⁷ used ANNs to derive insight into the memorability of images in humans. Nayebi et al.¹⁸⁵ studied the impact of combining embodiment and unsupervised learning on learning to explain the mouse visual cortex. What these examples have in common is that they use neuroconnectionist models to derive computational insights that may have been intractable without ANN simulations, due to the complexity of distributed neural codes.

These developments in the belt have raised new challenges in how we compare models and interpret results, requiring improvements in methodology. This has led to developments in quantifying the alignment of ANN and brain data^{157,158,163,198,199}, training networks end-to-end directly to match neural data^{19,77,92–99}, highlighting the importance individual variability across network instances¹⁶⁷, explicitly pitting networks against each other using “controversial stimuli”²⁰⁰, integrating hierarchical⁷⁴ and temporal^{19,150} aspects of information processing into model comparisons, performing detailed psychophysical experiments^{140,144,150,201}, using metamers to compare humans and models²⁰², and developing interpretable low-rank recurrent networks²⁰³.

Altogether, these rapid developments of neuroconnectionism demonstrate that the underlying research programme is highly progressive. Researchers have repeatedly updated their models, altering the architectures, training objectives, and datasets to arrive at a progressively better account of neural information processing, including ANN-driven new insights and hypotheses. Although we are still far from final answers, and there are still important methodological issues^{158,204–206}, this illustrates how neuroconnectionist experiments, theory and methodology form a virtuous circle in which new empirical and theoretical results require improved methodology, which in turn allows for new results, allowing the field to keep progressing.

Shortcomings as Signposts, not Roadblocks.

At the same time, ANNs as models of brain function continue to be met with skepticism as differences with brain data and behavior remain. In the following, we consider prominent examples, asking whether the underlying controversies are *roadblocks* that preclude progress (and thereby may turn neuroconnectionism into a degenerating research programme), or whether they are better perceived as important *signposts*, which are pointing towards promising directions of improvement. In visiting these controversies, we highlight recent developments in the field that aim to solve them, suggesting that neuroconnectionism is a progressive research programme.

One of the main controversies surrounding ANNs in both cognitive science and artificial intelligence, concerns differences in the behavior of ANNs vs. biological brains in a variety of settings^{53,56,207}. Although capable of impressive behavior, **current ANNs do not model all aspects of behavior equally well: they generalize poorly¹⁴², can be easily fooled²⁰⁸⁻²¹⁰ and behave differently from humans in many psychophysical settings^{141,144,211}**. These observations are important for the neuroconnectionist programme, as they point out instances where the current set of models exhibit shortcomings. Researchers are exploring a variety of avenues to address these shortcomings and improve the biological match of ANNs. As detailed above, attention mechanisms²¹²⁻²¹⁴, recurrent mechanisms^{19,76,80,82,139,174}, foveation²¹⁵, and biology-inspired receptive fields in early layers²¹⁶⁻²¹⁹ have been shown to improve different aspects of behavioral alignment with biology. Hence, the current shortcomings of ANNs in reproducing human behavior are not hindering the progress of neuroconnectionism. Rather, they are serving as inspiration for a multitude of new developments. They are real challenges that need to be taken seriously, but they are signposts, not roadblocks.

Another controversy focuses on computations on the level of symbolic manipulations in ANN vs biology. It has been argued that current ANNs, unlike humans, do not compute with symbols because they do not have appropriately structured representations, systematicity, dynamic variable binding and role-filler independence, which may ultimately cause tractability/scalability issues^{129,220}. The field of neurosymbolic computations is a response to such criticisms²²¹⁻²²⁴. Moreover, even the most ardent defenders of symbolism agree that ANNs should in principle be able to implement the symbolic systems we all desire^{129,220,225}. Indeed, such a capacity is an immediate result of the universal approximation theorems that hold for neural networks^{67,226}. Again, symbolic computations are not a roadblock for neuroconnectionism, they are signposts pointing to important directions of research that can be addressed from within the research program. The same goes for causal modeling²²⁷ and related arguments.

Discussions surrounding ANNs as modeling framework are also based on their level of complexity. On the one hand, it is reasoned that (current) **ANNs are too abstract, and that omitting many biological details renders them unable to adequately model neural processes. Indeed, current ANNs are vastly simpler than the breathtaking complexity of biological neural networks**, which involves, to name a few examples, intricate dendritic computations^{228,229}, neurotransmitter dynamics, and communication via temporally fine-grained spikes^{86,87,230}. At the same time, others deem ANNs *too complex* to provide a deeper understanding of cognition^{54,231}. Indeed, the millions of parameters of ANNs constitute an important departure from classic models with only few interpretable parameters or experimental conditions. Critics ask, what is gained by modeling one complex system with another?

So are ANNs too simple or too complex? As mentioned in the introduction, although current ANNs are likely too simple, it is an empirical question at what level of detail cognitive systems are best understood. Instead of committing to a single level, the neuroconnectionist research programme is capable of modeling different levels of abstraction, which allows researchers to implement biological detail in a top-down fashion through the testing of hypotheses about which biological features are important for modeling cognitive phenomena^{46,232}. **We currently do not know which biological features are central for brain computations, and which are merely byproducts of evolution that could be equivalently implemented in other ways.** Slowly testing the impact of more biological features in a hypothesis-driven fashion provides a way to simultaneously improve current ANNs and find out

which biological features are important. At the same time, the comparably large number of ANN parameters is likely necessary to encode important domain knowledge, and to perform computations end-to-end from raw sensory input to behavioral output. Relatedly, the complexity of a model needs to be evaluated in relation to the complexity of the system being modeled. For example, the number of units in Alexnet (659k), a commonly used model in visual computational neuroscience, is equivalent to the number of neurons in a cube of cortex with an edge length of 2.3mm. Furthermore, the number of parameters is a poor proxy of model complexity in ANNs, and the degrees of freedom manipulated by the experimenter (i.e. architectural and learning hyperparameters) are orders of magnitude smaller.

In summary, current ANNs are both simple (e.g. in terms of architectural complexity), and complex (e.g. in terms of the number of adjustable parameters). But, to proponents of neuroconnectionism, this very balance is a feature and not a bug, since it allows for top-down testing of which biological details matter. As reviewed above, ANNs are actively moving towards increased realism in terms of architectures used, input datasets, as well as training objectives. The integration of biological detail, however, is driven as part of a hypothesis testing approach. This is importantly different from mirroring biology for its own sake, because biological details come at the cost of computational complexity and explanatory merit, giving good reason to try and discard the less relevant biological aspects.

Crucially, determining what level of complexity our models require in order to understand higher-level cognitive phenomena is itself an empirical project. Neuroconnectionism offers an ideal framework for studying which biological details are necessary and which are not, while at the same time being expressive enough to perform complex tasks. Put simply, *neuroconnectionist models live in a Goldilocks Zone between computational abstraction and biological detail*. They are sufficiently abstract to make them tractable and trainable (given the currently available methods and compute), but retain sufficient biological detail in their algorithmic structure to map them onto neural and behavioral data.

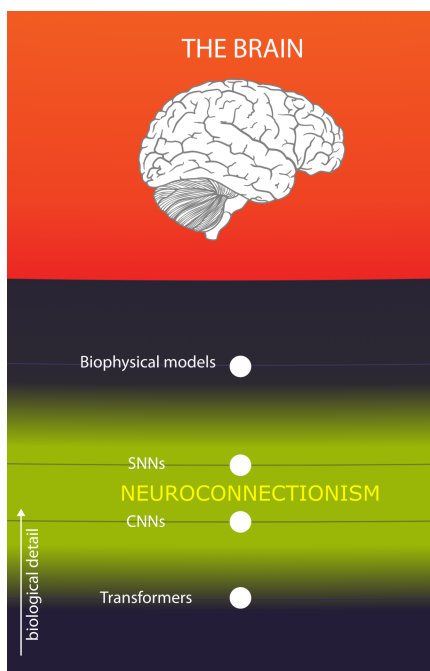


Figure 4: The Goldilocks zone of biological abstraction. Neuroconnectionism offers a coherent and computationally tractable framework to model cognition in which models vary in how much they abstract away from biology. This enables researchers to determine which biological details are needed to model essential aspects of brain computations. Hence, neuroconnectionism provides a goldilocks zone of biological abstraction, abstract enough to build large models capable of performing cognitive tasks, yet biologically detailed enough to test specific neuroscientific hypotheses.

Conclusions

In less than a decade, ANN modeling has gone from being fringe to a more central research tool in many parts of cognitive neuroscience, including vision, audition, and on to language and higher-level cognitive tasks. The underlying neuroconnectionism research programme has met both striking successes but also challenges. How should these contrasting results be interpreted? We have argued that the framework of progressive vs. degenerating Lakatosian research programmes is well suited to address this question. **Of central importance is the view that neuroconnectionism is not a single theory or hypothesis. Instead, it is a research programme composed of many different auxiliary hypotheses and research directions in the belt, each sharing the same core approach to model cognitive phenomena through distributed neural communications in ANNs.** Based on this view, we showed that current challenges represent important signposts that aid further progress, rather than roadblocks. Indeed, many such challenges have sparked vibrant new research directions. Together with the growing body of studies demonstrating good agreement between connectionist systems and brain data, these observations illustrate that the neuroconnectionism research programme is highly progressive. Importantly, traditional approaches based on simple human-interpretable concepts face significant challenges that can be overcome by the addition of neuroconnectionist models to the researcher's toolbox. These include explanations of neural data collected under more natural and varied stimulus paradigms, explanations of lower-level sensory processes and their transition to higher-level conceptual information, explanations of behavioral effects based on raw sensory input, and learning. Another argument supporting the use of ANNs in neuroscience is that theoretical computer science has taught us that **complex systems often cannot be simplified to predict their behavior: the only way to know the outcome of processing is to actually run the system** (or a good enough ANN model of it).

One reason for the success of ANNs may be that they provide a useful level of granularity for cognitive neuroscience. They provide a sufficiently abstract view on neural processes while being able to incrementally test which biological detail is needed in a top-down, hypothesis-driven manner. The hyperparameters of the training procedure thereby form the effective degrees of freedom of the experimenter, not the number of parameters in the respective network. On the other hand, ANNs are high-dimensional enough to encode domain knowledge, enabling them to be grounded in sensory input while performing comparatively well on a set of behavioral tasks.

While we have here focused on the progressive nature of the research programme, we do not mean to imply that the field is anywhere near to successfully explaining cognition. Neuroconnectionism is still in its infancy, and hence knowing where it fails is equally important as knowing where it works. In addition to addressing current challenges, future work will need to incorporate many currently missing aspects of lower and higher cognition. Focusing purely on the case of visual processing for the sake of argument, these include (i) multi-task networks that can explain our ability to perform multiple, sometimes highly abstract tasks based on sensory evidence²³³, (ii) a focus on embodiment rather than treating networks as artificial brains in vats^{179,185,186,234}, (iii) data efficient and continuous learning using biologically realistic learning rules and inductive biases, (iv) an answer to how symbolic reasoning is implemented in neuroconnectionist models, (v) better methods for unsupervised, multimodal learning, (vi) better modeling of cognitive development, (vii) the integration of multiple memory systems, and (viii) learning in social context (learning from each other). The fact that all these aspects are currently missing but can in principle be implemented in ANNs perfectly illustrates

both the current limitations and strong potential of the neuroconnectionist approach. Rather than feeling threatened by the possibility of a new connectionist winter, the neuroconnectionist community should therefore continue to welcome criticism and limitations as they point the way towards new insights. Critics of neuroconnectionism on the other hand should not regard every shortcoming of the current set of networks as a failure of the entire research programme. Time will tell whether neuroconnectionism can deliver on its promises to explain the emergence of cognitive phenomena, behavior and neural data from bio-inspired, yet simple distributed coding principles. For now, it remains a highly progressive and therefore exciting research programme that welcomes critical signposts to guide the way.

Acknowledgements

The authors acknowledge support by SNF grant n.203018 (Doerig), the ERC stg grant 101039524 TIME (Kietzmann), and the Max Planck Research Group grant of Martin N. Hebart (Seeliger).

References

1. Churchland, P. S. & Sejnowski, T. J. Blending computational and experimental neuroscience. *Nature Reviews Neuroscience* **17**, 667–668 (2016).
2. Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., Maclver, M. A. & Poeppel, D. Neuroscience needs behavior: correcting a reductionist bias. *Neuron* **93**, 480–490 (2017).
3. Kanwisher, N., McDermott, J. & Chun, M. M. The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *J. Neurosci.* **17**, 4302–4311 (1997).
4. Sergent, J., Ohta, S. & Macdonald, B. Functional neuroanatomy of face and object processing: a positron emission tomography study. *Brain* **115**, 15–36 (1992).
5. Tong, F., Nakayama, K., Vaughan, J. T. & Kanwisher, N. Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron* **21**, 753–759 (1998).
6. Tsao, D. Y., Freiwald, W. A., Knutsen, T. A., Mandeville, J. B. & Tootell, R. B. Faces and objects in macaque cerebral cortex. *Nature neuroscience* **6**, 989–995 (2003).
7. Rust, N. C. & Movshon, J. A. In praise of artifice. *Nature neuroscience* **8**, 1647–1650 (2005).
8. Vinken, K., Konkle, T. & Livingstone, M. The neural code for ‘face cells’ is not face specific. *bioRxiv* (2022).
9. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
10. Schmidhuber, J. Deep learning in neural networks: An overview. *Neural networks* **61**, 85–117 (2015).
11. McCulloch, W. S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* **5**, 115–133 (1943).
12. Cichy, R. M. & Kaiser, D. Deep neural networks as scientific models. *Trends in Cognitive Sciences* **23**, 305–317 (2019).
13. Schrimpf, M. *et al.* Brain-Score: Which Artificial Neural Network for Object Recognition is most Brain-Like? *bioRxiv* (2018) doi:10.1101/407007.
14. Rumelhart, D. E., McClelland, J. L. & Group, P. R. *Parallel distributed processing*. vol. 1 (IEEE New York, 1988).
15. Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A. & Oliva, A. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific reports* **6**, 1–13 (2016).
16. Fukushima, K. & Miyake, S. Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. in *Competition and cooperation in neural nets*

- 267–285 (Springer, 1982).
17. Güçlü, U. & Gerven, M. A. J. van. Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream. *J. Neurosci.* **35**, 10005–10014 (2015).
 18. Khaligh-Razavi, S.-M. & Kriegeskorte, N. Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS computational biology* **10**, e1003915 (2014).
 19. Kietzmann, T. C. *et al.* Recurrence is required to capture the representational dynamics of the human visual system. *Proceedings of the National Academy of Sciences* **116**, 21854–21863 (2019).
 20. Seeliger, K. *et al.* Convolutional neural network-based encoding and decoding of visual object recognition in space and time. *NeuroImage* **180**, 253–266 (2018).
 21. Yamins, D. L. *et al.* Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences* **111**, 8619–8624 (2014).
 22. Kell, A. J., Yamins, D. L., Shook, E. N., Norman-Haignere, S. V. & McDermott, J. H. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron* **98**, 630–644. e16 (2018).
 23. Saddler, M. R., Gonzalez, R. & McDermott, J. H. Deep neural network models reveal interplay of peripheral coding and stimulus statistics in pitch perception. *Nature communications* **12**, 1–25 (2021).
 24. Cadena, S. A. *et al.* Diverse task-driven modeling of macaque V4 reveals functional specialization towards semantic tasks. *bioRxiv* (2022).
 25. Jackson, R. L., Rogers, T. T. & Lambon Ralph, M. A. Reverse-engineering the cortical architecture for controlled semantic cognition. *Nature human behaviour* **5**, 774–786 (2021).
 26. Saxe, A. M., McClelland, J. L. & Ganguli, S. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences* **116**, 11537–11546 (2019).
 27. Caucheteux, C. & King, J.-R. Brains and algorithms partially converge in natural language processing. *Communications biology* **5**, 1–10 (2022).
 28. Schrimpf, M. *et al.* Artificial neural networks accurately predict language processing in the brain. *BioRxiv* **10**, 26.174482 (2020).
 29. Hannagan, T., Agrawal, A., Cohen, L. & Dehaene, S. Emergence of a compositional neural code for written words: Recycling of a convolutional neural network for reading. *Proceedings of the National Academy of Sciences* **118**, (2021).
 30. Botvinick, M., Wang, J. X., Dabney, W., Miller, K. J. & Kurth-Nelson, Z. Deep reinforcement learning and its neuroscientific implications. *Neuron* (2020).
 31. Dabney, W. *et al.* A distributional code for value in dopamine-based reinforcement learning. *Nature* **577**, 671–675 (2020).
 32. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature* **503**, 78–84 (2013).
 33. Quax, S. & van Gerven, M. Emergent mechanisms of evidence integration in recurrent neural networks. *PLoS One* **13**, e0205676 (2018).
 34. Orhan, A. E. & Ma, W. J. A diverse range of factors affect the nature of neural representations underlying short-term memory. *Nature neuroscience* **22**, 275–283 (2019).
 35. Cross, L., Cockburn, J., Yue, Y. & O’Doherty, J. P. Using deep reinforcement learning to reveal how the brain encodes abstract state-space representations in high-dimensional environments. *Neuron* **109**, 724–738. e7 (2021).
 36. Feulner, B. *et al.* Small, correlated changes in synaptic connectivity may facilitate rapid motor learning. *Nature Communications* **13**, 1–14 (2022).
 37. Merel, J., Botvinick, M. & Wayne, G. Hierarchical motor control in mammals and machines. *Nature communications* **10**, 1–12 (2019).
 38. Michaels, J. A., Schaffelhofer, S., Agudelo-Toro, A. & Scherberger, H. A goal-driven modular neural network predicts parietofrontal neural dynamics during grasping. *Proceedings of the*

- national academy of sciences* **117**, 32124–32135 (2020).
39. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature neuroscience* **18**, 1025–1033 (2015).
 40. Bao, P., She, L., McGill, M. & Tsao, D. Y. A map of object space in primate inferotemporal cortex. *Nature* **583**, 103–108 (2020).
 41. Blauch, N. M., Behrmann, M. & Plaut, D. C. A connectivity-constrained computational account of topographic organization in primate high-level visual cortex. *Proceedings of the National Academy of Sciences* **119**, (2022).
 42. Dobs, K., Martinez, J., Kell, A. J. E. & Kanwisher, N. Brain-like functional specialization emerges spontaneously in deep neural networks. *Science Advances* **8**, eabl8913 (2022).
 43. Doerig, A., Krahmer, B. & Kietzmann, T. Emergence of topographic organization in a non-convolutional deep neural network (Neuromatch 40). in (2021).
 44. Higgins, I. *et al.* Unsupervised deep learning identifies semantic disentanglement in single inferotemporal face patch neurons. *Nature communications* **12**, 1–14 (2021).
 45. Lee, H. *et al.* Topographic deep artificial neural networks reproduce the hallmarks of the primate inferior temporal cortex face processing network. *bioRxiv* (2020).
 46. Kietzmann, T. C., McClure, P. & Kriegeskorte, N. Deep neural networks in computational neuroscience. in *Oxford research encyclopedia of neuroscience* (2019).
 47. Kriegeskorte, N. Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science* **1**, 417–446 (2015).
 48. Lindsay, G. W. Convolutional neural networks as a model of the visual system: Past, present, and future. *Journal of cognitive neuroscience* **33**, 2017–2031 (2021).
 49. Marblestone, A. H., Wayne, G. & Kording, K. P. Toward an integration of deep learning and neuroscience. *Frontiers in computational neuroscience* **94** (2016).
 50. Richards, B. A. *et al.* A deep learning framework for neuroscience. *Nat Neurosci* **22**, 1761–1770 (2019).
 51. Saxe, A., Nelli, S. & Summerfield, C. If deep learning is the answer, what is the question? *Nature Reviews Neuroscience* 1–13 (2020).
 52. Van Gerven, M. Computational foundations of natural intelligence. *Frontiers in computational neuroscience* **112** (2017).
 53. Bowers, J. S. *et al.* Deep Problems with Neural Network Models of Human Vision. (2022).
 54. Leek, E. C., Leonardis, A. & Heinke, D. Deep neural networks and image classification in biological vision. *Vision Research* **197**, 108058 (2022).
 55. Marcus, G. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631* (2018).
 56. Serre, T. Deep learning: the good, the bad, and the ugly. *Annual Review of Vision Science* **5**, 399–426 (2019).
 57. Lakatos, I. Falsification and the methodology of scientific research programmes. in *Can theories be refuted?* 205–259 (Springer, 1976).
 58. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. in *Advances in neural information processing systems* 1097–1105 (2012).
 59. LeCun, Y. *et al.* Handwritten digit recognition with a back-propagation network. in *Advances in neural information processing systems* 396–404 (1990).
 60. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology* **160**, 106–154 (1962).
 61. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. in *Proceedings of the IEEE conference on computer vision and pattern recognition* 770–778 (2016).
 62. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *nature* **323**, 533–536 (1986).
 63. Hebb, D. O. *The organization of behavior: A neuropsychological theory.* (Psychology Press, 2005).

64. Rao, R. P. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience* **2**, 79 (1999).
65. Cook, M. Universality in elementary cellular automata. *Complex systems* **15**, 1–40 (2004).
66. Wolfram, S. Cellular automata as models of complexity. *Nature* **311**, 419–424 (1984).
67. Siegelmann, H. T. & Sontag, E. D. On the computational power of neural nets. *Journal of computer and system sciences* **50**, 132–150 (1995).
68. Wittgenstein, L. *Philosophical investigations*. (John Wiley & Sons, 2009).
69. Ali, A., Ahmad, N., de Groot, E., van Gerven, M. A. J. & Kietzmann, T. C. Predictive coding is a consequence of energy efficiency in recurrent neural networks. *bioRxiv* 1–17 (2021) doi:10.1101/2021.02.16.430904.
70. Heilbron, M., Armeni, K., Schoffelen, J.-M., Hagoort, P. & De Lange, F. P. A hierarchy of linguistic predictions during natural language comprehension. *Proceedings of the National Academy of Sciences* **119**, e2201968119 (2022).
71. Ponce, C. R. *et al.* Evolving images for visual neurons using a deep generative network reveals coding principles and neuronal preferences. *Cell* **177**, 999–1009. e10 (2019).
72. Tuli, S., Dasgupta, I., Grant, E. & Griffiths, T. L. Are Convolutional Neural Networks or Transformers more like human vision? *arXiv preprint arXiv:2105.07197* (2021).
73. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
74. Nonaka, S., Majima, K., Aoki, S. C. & Kamitani, Y. Brain hierarchy score: Which deep neural networks are hierarchically brain-like? *Iscience* **24**, 103013 (2021).
75. Cichy, R. M. *et al.* The algonauts project: A platform for communication between the sciences of biological and artificial intelligence. *arXiv preprint arXiv:1905.05675* (2019).
76. Doerig, A., Schmittwilken, L., Sayim, B., Manassi, M. & Herzog, M. H. Capsule networks as recurrent models of grouping and segmentation. *PLoS computational biology* **16**, e1008017 (2020).
77. Güçlü, U. & Van Gerven, M. A. Modeling the dynamics of human brain activity with recurrent neural networks. *Frontiers in computational neuroscience* **11**, 7 (2017).
78. Kar, K. & DiCarlo, J. J. Fast recurrent processing via ventrolateral prefrontal cortex is needed by the primate ventral stream for robust core visual object recognition. *Neuron* **109**, 164–176. e5 (2021).
79. Lindsay, G. W., Mrcic-Flogel, T. D. & Sahani, M. Bio-inspired neural networks implement different recurrent visual processing strategies than task-trained ones do. *bioRxiv* (2022).
80. Linsley, D., Kim, J. & Serre, T. Sample-efficient image segmentation through recurrence. *arXiv:1811.11356 [cs]* (2018).
81. Nayebi, A. *et al.* Goal-Driven Recurrent Neural Network Models of the Ventral Visual Stream. *bioRxiv* (2021).
82. Thorat, S., Aldegheri, G. & Kietzmann, T. C. Category-orthogonal object features guide information processing in recurrent neural networks trained for object categorization. *arXiv preprint arXiv:2111.07898* (2021).
83. Bertalmío, M. *et al.* Evidence for the intrinsically nonlinear nature of receptive fields in vision. *Scientific reports* **10**, 1–15 (2020).
84. Quax, S. C., D’Asaro, M. & van Gerven, M. A. Adaptive time scales in recurrent neural networks. *Scientific reports* **10**, 1–14 (2020).
85. Voelker, A., Kajić, I. & Eliasmith, C. Legendre memory units: Continuous-time representation in recurrent neural networks. *Advances in neural information processing systems* **32**, (2019).
86. Bohte, S. M. The evidence for neural information processing with precise spike-times: A survey. *Natural Computing* **3**, 195–206 (2004).
87. Gerstner, W. & Kistler, W. M. *Spiking neuron models: Single neurons, populations, plasticity*. (Cambridge university press, 2002).
88. Sörensen, L. K., Zambrano, D., Slagter, H. A., Bohté, S. M. & Scholte, H. S. Leveraging spiking

- deep neural networks to understand the neural mechanisms underlying selective attention. *Journal of cognitive neuroscience* **34**, 655–674 (2022).
89. Zenke, F. & Ganguli, S. Superspike: Supervised learning in multilayer spiking neural networks. *Neural computation* **30**, 1514–1541 (2018).
 90. Guerguiev, J., Lillicrap, T. P. & Richards, B. A. Towards deep learning with segregated dendrites. *ELife* **6**, e22901 (2017).
 91. Sacramento, J., Ponte Costa, R., Bengio, Y. & Senn, W. Dendritic cortical microcircuits approximate the backpropagation algorithm. *Advances in neural information processing systems* **31**, (2018).
 92. Antolík, J., Hofer, S. B., Bednar, J. A. & Mrcic-Flogel, T. D. Model constrained by visual hierarchy improves prediction of neural responses to natural scenes. *PLoS computational biology* **12**, e1004927 (2016).
 93. Cadena, S. A. *et al.* Deep convolutional models improve predictions of macaque V1 responses to natural images. *PLoS computational biology* **15**, e1006897 (2019).
 94. Ecker, A. S. *et al.* A rotation-equivariant convolutional neural network model of primary visual cortex. *arXiv preprint arXiv:1809.10504* (2018).
 95. Kindel, W. F., Christensen, E. D. & Zylberberg, J. Using deep learning to probe the neural code for images in primary visual cortex. *Journal of vision* **19**, 29–29 (2019).
 96. Klindt, D. A., Ecker, A. S., Euler, T. & Bethge, M. Neural system identification for large populations separating "what" and "where". *arXiv preprint arXiv:1711.02653* (2017).
 97. Seeliger, K. *et al.* End-to-end neural system identification with neural information flow. *PLOS Computational Biology* **17**, e1008558 (2021).
 98. St-Yves, G. & Naselaris, T. The feature-weighted receptive field: an interpretable encoding model for complex feature spaces. *NeuroImage* **180**, 188–202 (2018).
 99. Tripp, B. Approximating the architecture of visual cortex in a convolutional network. *Neural computation* **31**, 1551–1591 (2019).
 100. Allen, E. J. *et al.* A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience* **25**, 116–126 (2022).
 101. Bellec, P. & Boyle, J. Bridging the gap between perception and action: the case for neuroimaging, AI and video games. (2019).
 102. Hebart, M. N. *et al.* THINGS: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLoS one* **14**, e0223792 (2019).
 103. Naselaris, T., Allen, E. & Kay, K. Extensive sampling for complete models of individual brains. *Current Opinion in Behavioral Sciences* **40**, 45–51 (2021).
 104. Seeliger, K., Sommers, R. P., Güçlü, U., Bosch, S. E. & Van Gerven, M. A. J. A large single-participant fMRI dataset for probing brain responses to naturalistic stimuli in space and time. *bioRxiv* 687681 (2019).
 105. Siegle, J. H. *et al.* Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).
 106. Mehrer, J., Spoerer, C. J., Jones, E. C., Kriegeskorte, N. & Kietzmann, T. C. An ecologically motivated image dataset for deep learning yields better models of human vision. *Proceedings of the National Academy of Sciences* **118**, e2011417118 (2021).
 107. Chen, T., Kornblith, S., Norouzi, M. & Hinton, G. A simple framework for contrastive learning of visual representations. in *International conference on machine learning* 1597–1607 (PMLR, 2020).
 108. Konkle, T. & Alvarez, G. A. Instance-level contrastive learning yields human brain-like representation without category-supervision. *BioRxiv* 2020.06. 15.153247 (2020).
 109. Choksi, B. *et al.* Predify: Augmenting deep neural networks with brain-inspired predictive coding dynamics. *Advances in Neural Information Processing Systems* **34**, 14069–14083 (2021).
 110. Lotter, W., Kreiman, G. & Cox, D. A neural network trained for prediction mimics diverse features of biological neurons and perception. *Nature Machine Intelligence* **2**, 210–219 (2020).

111. Soulos, P. & Isik, L. Disentangled face representations in deep generative models and the human brain. in *NeurIPS 2020 Workshop SVRHM* (2020).
112. Storrs, K. R., Anderson, B. L. & Fleming, R. W. Unsupervised learning predicts human perception and misperception of gloss. *Nature Human Behaviour* **5**, 1402–1417 (2021).
113. Franzius, M., Sprekeler, H. & Wiskott, L. Slowness and sparseness lead to place, head-direction, and spatial-view cells. *PLoS computational biology* **3**, e166 (2007).
114. Franzius, M., Wilbert, N. & Wiskott, L. Invariant object recognition with slow feature analysis. in *International Conference on Artificial Neural Networks* 961–970 (Springer, 2008).
115. Kayser, C., Einhäuser, W., Dümmer, O., König, P. & Körding, K. Extracting slow subspaces from natural videos leads to complex cells. in *International conference on artificial neural networks* 1075–1080 (Springer, 2001).
116. Wiskott, L. & Sejnowski, T. J. Slow feature analysis: Unsupervised learning of invariances. *Neural computation* **14**, 715–770 (2002).
117. Wyss, R., König, P. & Verschure, P. F. J. A model of the ventral visual system based on temporal stability and local memory. *PLoS biology* **4**, e120 (2006).
118. Lindsay, G. W., Merel, J., Mrsic-Flogel, T. & Sahani, M. Divergent representations of ethological visual inputs emerge from supervised, unsupervised, and reinforcement learning. *arXiv preprint arXiv:2112.02027* (2021).
119. Dwivedi, K., Bonner, M. F., Cichy, R. M. & Roig, G. Unveiling functions of the visual cortex using task-specific deep neural networks. *PLoS computational biology* **17**, e1009267 (2021).
120. Jang, H., McCormack, D. & Tong, F. Noise-trained deep neural networks effectively predict human vision and its neural responses to challenging images. *PLoS Biology* **19**, e3001418 (2021).
121. Ahmad, N., Schrader, E. & van Gerven, M. Constrained Parameter Inference as a Principle for Learning. *arXiv preprint arXiv:2203.13203* (2022).
122. Lillicrap, T. P., Cownden, D., Tweed, D. B. & Akerman, C. J. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications* **7**, 1–10 (2016).
123. Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J. & Hinton, G. Backpropagation and the brain. *Nature Reviews Neuroscience* **21**, 335–346 (2020).
124. Pozzi, I., Bohte, S. & Roelfsema, P. Attention-Gated Brain Propagation: How the brain can implement reward-based error backpropagation. *Advances in Neural Information Processing Systems* **33**, 2516–2526 (2020).
125. Richards, B. A. & Lillicrap, T. P. Dendritic solutions to the credit assignment problem. *Current opinion in neurobiology* **54**, 28–36 (2019).
126. Saxe, A. M., McClelland, J. L. & Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120* (2013).
127. Benjamin, A. S., Zhang, L.-Q., Qiu, C., Stocker, A. & Körding, K. P. Efficient neural codes naturally emerge through gradient descent learning. 2022.05.11.491548 Preprint at <https://doi.org/10.1101/2022.05.11.491548> (2022).
128. Munakata, Y. & Pfaffly, J. Hebbian learning and development. *Developmental science* **7**, 141–148 (2004).
129. Fodor, J. A. & Pylyshyn, Z. W. Connectionism and cognitive architecture: A critical analysis. *Cognition* **28**, 3–71 (1988).
130. Berrios, W. & Deza, A. Joint rotational invariance and adversarial training of a dual-stream Transformer yields state of the art Brain-Score for Area V4. *arXiv preprint arXiv:2203.06649* (2022).
131. St-Yves, G., Allen, E. J., Wu, Y., Kay, K. & Naselaris, T. Brain-optimized neural networks learn non-hierarchical models of representation in human visual cortex. *bioRxiv* (2022).
132. Hasenstaub, A., Otte, S., Callaway, E. & Sejnowski, T. J. Metabolic cost as a unifying principle governing neuronal biophysics. *Proceedings of the National Academy of Sciences* **107**, 12329–12334 (2010).

133. Stone, J. V. *Principles of neural information theory: computational neuroscience and metabolic efficiency (tutorial introductions)*. (Tutorial Introductions, 2018).
134. Wang, Z., Wei, X.-X., Stocker, A. A. & Lee, D. D. Efficient neural codes under metabolic constraints. *Advances in neural information processing systems* **29**, (2016).
135. Dosovitskiy, A. *et al.* An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
136. Silver, D. *et al.* Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
137. Mnih, V. *et al.* Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
138. Vinyals, O. *et al.* Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575**, 350–354 (2019).
139. Spoerer, C. J., Kietzmann, T. C., Mehrer, J., Charest, I. & Kriegeskorte, N. Recurrent neural networks can explain flexible trading of speed and accuracy in biological vision. *PLoS computational biology* **16**, e1008215 (2020).
140. Geirhos, R., Meding, K. & Wichmann, F. A. Beyond accuracy: quantifying trial-by-trial behaviour of CNNs and humans by measuring error consistency. *Advances in Neural Information Processing Systems* **33**, 13890–13902 (2020).
141. Geirhos, R. *et al.* ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. in *International Conference on Learning Representations* (2018).
142. Geirhos, R. *et al.* Generalisation in humans and deep neural networks. *Advances in neural information processing systems* **31**, (2018).
143. Singer, J. J., Seeliger, K., Kietzmann, T. C. & Hebart, M. N. From photos to sketches-how humans and deep neural networks process objects across different levels of visual abstraction. *Journal of vision* **22**, 4–4 (2022).
144. Doerig, A., Bornet, A., Choung, O. H. & Herzog, M. H. Crowding reveals fundamental differences in local vs. global processing in humans and machines. *Vision Research* **167**, 39–45 (2020).
145. Funke, C. M. *et al.* Comparing the ability of humans and DNNs to recognise closed contours in cluttered images. in *18th Annual Meeting of the Vision Sciences Society (VSS 2018)* 213 (2018).
146. Jacob, G., Pramod, R. T., Katti, H. & Arun, S. P. Qualitative similarities and differences in visual object representations between brains and deep networks. *Nature communications* **12**, 1–14 (2021).
147. Kim, J., Linsley, D., Thakkar, K. & Serre, T. Disentangling neural mechanisms for perceptual grouping. *arXiv preprint arXiv:1906.01558* (2019).
148. Loke, J. *et al.* A critical test of deep convolutional neural networks' ability to capture recurrent processing in the brain using visual masking. *bioRxiv* (2022).
149. RichardWebster, B., Anthony, S. & Scheirer, W. Psyphy: A psychophysics driven evaluation framework for visual recognition. *IEEE transactions on pattern analysis and machine intelligence* (2018).
150. Sörensen, L. K., Bohté, S. M., De Jong, D., Slagter, H. A. & Scholte, H. S. Mechanisms of human dynamic object recognition revealed by sequential deep neural networks. *bioRxiv* (2022).
151. Firestone, C. Performance vs. competence in human–machine comparisons. *Proceedings of the National Academy of Sciences* **117**, 26562–26571 (2020).
152. Lonnqvist, B., Bornet, A., Doerig, A. & Herzog, M. H. A comparative biology approach to DNN modeling of vision: A focus on differences, not similarities. *Journal of vision* **21**, 17–17 (2021).
153. Ma, W. J. & Peters, B. A neural network walks into a lab: towards using deep nets as models for human behavior. *arXiv preprint arXiv:2005.02181* (2020).
154. Neri, P. Deep networks may capture biological behavior for shallow, but not deep, empirical characterizations. *Neural Networks* **152**, 244–266 (2022).
155. Kriegeskorte, N., Mur, M. & Bandettini, P. A. Representational similarity analysis-connecting the

- branches of systems neuroscience. *Frontiers in systems neuroscience* **2**, 4 (2008).
156. Kriegeskorte, N. & Wei, X.-X. Neural tuning and representational geometry. *Nature Reviews Neuroscience* **22**, 703–718 (2021).
 157. Kaniuth, P. & Hebart, M. N. Feature-reweighted representational similarity analysis: A method for improving the fit between computational models, brains, and behavior. *NeuroImage* 119294 (2022).
 158. Storrs, K. R., Kietzmann, T. C., Walther, A., Mehrer, J. & Kriegeskorte, N. Diverse deep neural networks all predict human inferior temporal cortex well, after training and fitting. *Journal of Cognitive Neuroscience* **33**, 2044–2064 (2021).
 159. Kornblith, S., Norouzi, M., Lee, H. & Hinton, G. Similarity of neural network representations revisited. in *International Conference on Machine Learning* 3519–3529 (PMLR, 2019).
 160. Kriegeskorte, N. & Diedrichsen, J. Peeling the onion of brain representations. *Annual review of neuroscience* **42**, 407–432 (2019).
 161. Naselaris, T., Kay, K. N., Nishimoto, S. & Gallant, J. L. Encoding and decoding in fMRI. *Neuroimage* **56**, 400–410 (2011).
 162. van Gerven, M. A. J. A primer on encoding models in sensory neuroscience. *Journal of Mathematical Psychology* **76**, 172–183 (2017).
 163. Sexton, N. J. & Love, B. C. Reassessing hierarchical correspondences between brain and deep networks through direct interface. *Science Advances* **8**, eabm2219 (2022).
 164. Bashivan, P., Kar, K. & DiCarlo, J. J. Neural population control via deep image synthesis. *Science* **364**, (2019).
 165. Gu, Z. *et al.* NeuroGen: activation optimized image synthesis for discovery neuroscience. *NeuroImage* **247**, 118812 (2022).
 166. Ratan Murty, N. A., Bashivan, P., Abate, A., DiCarlo, J. J. & Kanwisher, N. Computational models of category-selective brain regions enable high-throughput tests of selectivity. *Nature Communications* **12**, 1–14 (2021).
 167. Mehrer, J., Spoerer, C. J., Kriegeskorte, N. & Kietzmann, T. C. Individual differences among deep neural network models. *Nature Communications* **11**, 5725 (2020).
 168. Geadah, V., Horoi, S., Kerg, G., Wolf, G. & Lajoie, G. Goal-driven optimization of single-neuron properties in artificial networks reveals regularization role of neural diversity and adaptation. *bioRxiv* (2022).
 169. Doshi, F. & Konkle, T. Organizational motifs of cortical responses to objects emerge in topographic projections of deep neural networks. *Journal of Vision* **21**, 2226–2226 (2021).
 170. Elsayed, G., Ramachandran, P., Shlens, J. & Kornblith, S. Revisiting spatial invariance with low-rank local connectivity. in *International Conference on Machine Learning* 2868–2879 (PMLR, 2020).
 171. Zaadnoordijk, L., Besold, T. R. & Cusack, R. Lessons from infant learning for unsupervised machine learning. *Nature Machine Intelligence* **4**, 510–520 (2022).
 172. Zador, A. M. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature communications* **10**, 1–7 (2019).
 173. Rane, S. *et al.* Predicting Word Learning in Children from the Performance of Computer Vision Systems. Preprint at <https://doi.org/10.48550/arXiv.2207.09847> (2022).
 174. Kar, K., Kubilius, J., Schmidt, K., Issa, E. B. & DiCarlo, J. J. Evidence that recurrent circuits are critical to the ventral stream’s execution of core object recognition behavior. *Nature neuroscience* **22**, 974 (2019).
 175. Kreiman, G. & Serre, T. Beyond the feedforward sweep: feedback computations in the visual cortex. *Annals of the New York Academy of Sciences* **1464**, 222–241 (2020).
 176. Nayebi, A. *et al.* Recurrent Connections in the Primate Ventral Visual Stream Mediate a Trade-Off Between Task Performance and Network Size During Core Object Recognition. *Neural Computation* **34**, 1652–1675 (2022).
 177. van Bergen, R. S. & Kriegeskorte, N. Going in circles is the way forward: the role of recurrence in

- visual inference. *Current Opinion in Neurobiology* **65**, 176–193 (2020).
178. Sullivan, J., Mei, M., Perfors, A., Wojcik, E. & Frank, M. C. SAYCam: A large, longitudinal audiovisual dataset recorded from the infant’s perspective. *Open mind* **5**, 20–29 (2021).
 179. Clay, V., König, P., Kühnberger, K.-U. & Pipa, G. Learning sparse and meaningful representations through embodiment. *Neural Networks* **134**, 23–41 (2021).
 180. Gan, C. *et al.* The threedworld transport challenge: A visually guided task-and-motion planning benchmark for physically realistic embodied ai. *arXiv preprint arXiv:2103.14025* (2021).
 181. Chen, Y. *et al.* COCO-Search18 fixation dataset for predicting goal-directed attention control. *Scientific reports* **11**, 1–11 (2021).
 182. Konkle, T. & Alvarez, G. A. A self-supervised domain-general learning framework for human ventral stream representation. *Nature communications* **13**, 1–12 (2022).
 183. Zhuang, C. *et al.* Unsupervised neural network models of the ventral visual stream. *Proceedings of the National Academy of Sciences* **118**, (2021).
 184. Bakhtiari, S., Mineault, P., Lillicrap, T., Pack, C. & Richards, B. The functional specialization of visual cortex emerges from training parallel pathways with self-supervised predictive learning. *Advances in Neural Information Processing Systems* **34**, (2021).
 185. Nayebi, A. *et al.* Mouse visual cortex as a limited resource system that self-learns an ecologically-general representation. *bioRxiv* 2021.06.16.448730 (2022).
 186. Mineault, P., Bakhtiari, S., Richards, B. & Pack, C. Your head is there to move you around: Goal-driven models of the primate dorsal pathway. *Advances in Neural Information Processing Systems* **34**, (2021).
 187. Stringer, S. M., Rolls, E. T. & Trappenberg, T. P. Self-organizing continuous attractor network models of hippocampal spatial view cells. *Neurobiology of learning and memory* **83**, 79–92 (2005).
 188. Tsodyks, M. Attractor neural network models of spatial maps in hippocampus. *Hippocampus* **9**, 481–489 (1999).
 189. Uria, B. *et al.* The Spatial Memory Pipeline: a model of egocentric to allocentric understanding in mammalian brains. *BioRxiv* (2020).
 190. Whittington, J. C. *et al.* The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* **183**, 1249-1263. e23 (2020).
 191. Whittington, J. C., Warren, J. & Behrens, T. E. Relating transformers to models and neural representations of the hippocampal formation. *arXiv preprint arXiv:2112.04035* (2021).
 192. Riveland, R. & Pouget, A. A neural model of task compositionality with natural language instructions. *bioRxiv* (2022).
 193. Xu, P., Zhu, X. & Clifton, D. A. Multimodal Learning with Transformers: A Survey. *arXiv preprint arXiv:2206.06488* (2022).
 194. Seeliger, K., Roth, J., Schmid, T. & Hebart, M. Synthesizing preferred stimuli for individual voxels in the human visual system. *Journal of Vision* **21**, 2311–2311 (2021).
 195. Walker, E. Y. *et al.* Inception loops discover what excites neurons most using deep predictive models. *Nature neuroscience* **22**, 2060–2065 (2019).
 196. Averbek, B. B. Pruning recurrent neural networks replicates adolescent changes in working memory and reinforcement learning. *Proceedings of the National Academy of Sciences* **119**, e2121331119 (2022).
 197. Rust, N. C. & Jannuzi, B. G. Identifying Objects and Remembering Images: Insights From Deep Neural Networks. *Current Directions in Psychological Science* 09637214221083663 (2022).
 198. Ivanova, A. A. *et al.* Beyond linear regression: mapping models in cognitive neuroscience should align with research goals. *arXiv preprint arXiv:2208.10668* (2022).
 199. Peterson, J. C., Abbott, J. T. & Griffiths, T. L. Evaluating (and improving) the correspondence between deep neural networks and human representations. *Cognitive science* **42**, 2648–2669 (2018).

200. Golan, T., Raju, P. C. & Kriegeskorte, N. Controversial stimuli: Pitting neural networks against each other as models of human cognition. *Proceedings of the National Academy of Sciences* **117**, 29330–29337 (2020).
201. Biscione, V. & Bowers, J. S. Do dnns trained on natural images acquire gestalt properties? *arXiv preprint arXiv:2203.07302* (2022).
202. Feather, J., Durango, A., Gonzalez, R. & McDermott, J. Metamers of neural networks reveal divergence from human perceptual systems. *Advances in Neural Information Processing Systems* **32**, (2019).
203. Mastrogiuseppe, F. & Ostojic, S. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron* **99**, 609–623. e29 (2018).
204. Dujmović, M., Bowers, J., Adolfi, F. & Malhotra, G. The pitfalls of measuring representational similarity using representational similarity analysis. *bioRxiv* (2022).
205. Elmoznino, E. & Bonner, M. F. High-performing neural network models of visual cortex benefit from high latent dimensionality. *bioRxiv* (2022).
206. Schaeffer, R., Khona, M. & Fiete, I. R. No Free Lunch from Deep Learning in Neuroscience: A Case Study through Models of the Entorhinal-Hippocampal Circuit. in *ICML 2022 2nd AI for Science Workshop* (2022).
207. Crick, F. The recent excitement about neural networks. *Nature* **337**, 129–132 (1989).
208. Moosavi-Dezfooli, S.-M., Fawzi, A. & Frossard, P. Deepfool: a simple and accurate method to fool deep neural networks. in *Proceedings of the IEEE conference on computer vision and pattern recognition* 2574–2582 (2016).
209. Nguyen, A., Yosinski, J. & Clune, J. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. in *Proceedings of the IEEE conference on computer vision and pattern recognition* 427–436 (2015).
210. Szegedy, C. *et al.* Intriguing properties of neural networks. in *2nd International Conference on Learning Representations, ICLR 2014* (2014).
211. Baker, N., Lu, H., Erlikhman, G. & Kellman, P. J. Deep convolutional networks do not classify based on global object shape. *PLoS computational biology* **14**, e1006613 (2018).
212. Lindsay, G. W. & Miller, K. D. How biological attention mechanisms improve task performance in a large-scale visual system model. *ELife* **7**, e38105 (2018).
213. Vaswani, A. *et al.* Attention is all you need. *Advances in neural information processing systems* **30**, (2017).
214. Xu, K. *et al.* Show, attend and tell: Neural image caption generation with visual attention. in *International conference on machine learning* 2048–2057 (PMLR, 2015).
215. Deza, A. & Konkle, T. Emergent properties of foveated perceptual systems. *arXiv preprint arXiv:2006.07991* (2020).
216. Dapello, J. *et al.* Simulating a primary visual cortex at the front of CNNs improves robustness to image perturbations. *Advances in Neural Information Processing Systems* **33**, 13073–13087 (2020).
217. Harrington, A. & Deza, A. Finding Biological Plausibility for Adversarially Robust Features via Metameric Tasks. in *SVRHM 2021 Workshop@ NeurIPS* (2021).
218. Malhotra, G., Evans, B. D. & Bowers, J. S. Hiding a plane with a pixel: examining shape-bias in CNNs and the benefit of building in biological constraints. *Vision Research* **174**, 57–68 (2020).
219. Safarani, S. *et al.* Towards robust vision by multi-task learning on monkey visual cortex. *Advances in Neural Information Processing Systems* **34**, (2021).
220. Marcus, G. F. *The algebraic mind: Integrating connectionism and cognitive science*. (MIT press, 2003).
221. Garnelo, M., Arulkumaran, K. & Shanahan, M. Towards deep symbolic reinforcement learning. *arXiv preprint arXiv:1609.05518* (2016).
222. Graves, A., Wayne, G. & Danihelka, I. Neural Turing machines. *arXiv preprint arXiv:1410.5401* (2014).

223. Holyoak, K. J. The Proper Treatment of Symbols. *Cognitive dynamics: Conceptual and representational change in humans and machines* 229 (2000).
224. Smolensky, P., McCoy, R. T., Fernandez, R., Goldrick, M. & Gao, J. Neurocompositional computing: From the Central Paradox of Cognition to a new generation of AI systems. *arXiv preprint arXiv:2205.01128* (2022).
225. Pinker, S. & Prince, A. On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* **28**, 73–193 (1988).
226. Hornik, K., Stinchcombe, M. & White, H. Multilayer feedforward networks are universal approximators. *Neural networks* **2**, 359–366 (1989).
227. Pearl, J. & Mackenzie, D. *The book of why: the new science of cause and effect*. (Basic books, 2018).
228. Larkum, M. E., Wu, J., Duverdin, S. A. & Gidon, A. The guide to dendritic spikes of the mammalian cortex in vitro and in vivo. *Neuroscience* (2022).
229. London, M. & Häusser, M. Dendritic computation. *Annu. Rev. Neurosci.* **28**, 503–532 (2005).
230. Engel, A. K., König, P., Kreiter, A. K., Schillen, T. B. & Singer, W. Temporal coding in the visual cortex: new vistas on integration in the nervous system. *Trends in neurosciences* **15**, 218–226 (1992).
231. McCloskey, M. Networks and theories: The place of connectionism in cognitive science. *Psychological science* **2**, 387–395 (1991).
232. Kriegeskorte, N. & Douglas, P. K. Cognitive computational neuroscience. *Nature neuroscience* **21**, 1148–1160 (2018).
233. Flesch, T., Juechems, K., Dumbalska, T., Saxe, A. & Summerfield, C. Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* (2022).
234. Molano-Mazon, M. *et al.* NeuroGym: An open resource for developing and sharing neuroscience tasks. (2022).