Cognition of Time and Thinking Beyond

Zedong Bi

Abstract

A common research protocol in cognitive neuroscience is to train subjects to perform deliberately designed experiments while recording brain activity, with the aim of understanding the brain mechanisms underlying cognition. However, how the results of this protocol of research can be applied in technology is seldom discussed. Here, I review the studies on time processing of the brain as examples of this research protocol, as well as two main application areas of neuroscience (neuroengineering and brain-inspired artificial intelligence). Time processing is a fundamental dimension of cognition, and time is also an indispensable dimension of any real-world signal to be processed in technology. Therefore, one may expect that the studies of time processing in cognition profoundly influence brain-related technology. Surprisingly, I found that the results from cognitive studies on timing processing are hardly helpful in solving practical problems. This awkward situation may be due to the lack of generalizability of the results of cognitive studies, which are under well-controlled laboratory conditions, to real-life situations. This lack of generalizability may be rooted in the fundamental unknowability of the world (including cognition). Overall, this paper questions and criticizes the usefulness and prospect of the abovementioned research protocol of cognitive neuroscience. I then give three suggestions for future research. First, to improve the generalizability of research, it is better to study brain activity under real-life conditions instead of in well-controlled laboratory experiments. Second, to overcome the unknowability of the world, we can engineer an easily accessible surrogate of the object under investigation, so that we can predict the behavior of the object under investigation by experimenting on the surrogate. Third, the paper calls for technology-oriented research, with the aim of technology creation instead of knowledge discovery.

Keywords

Metascience · Unknowability · Cognitive neuroscience · Generalizability · Technology application

Z. Bi (✉)
Lingang Laboratory, Shanghai, China

Institute for Future, Qingdao University, Qingdao, China

School of Automation, Shandong Key Laboratory of Industrial Control Technology, Qingdao University, Qingdao, China
e-mail: bizedong@lglab.ac.cn

H. Merchant, V. de Lafuente (eds.), *Neurobiology of Interval Timing*, Advances in Experimental Medicine and Biology 1455, https://doi.org/10.1007/978-3-031-60183-5_10

Prologue

> Humans are great through reasoning, but are matured by recognizing the limitations of reasoning. ---Prof. Qing Liu, School of Politics and International Relations, East China Normal University, Shanghai

"So interesting! Why is it?" This is perhaps the question that inspires your curiosity about the brain and marks the beginning of a neuroscience research journey. However, before you embark on such an investigation, I advise to think twice about whether the mechanism underlying the phenomenon is worth investigating. In most cases, such investigation is useless in solving practical problems.

"Too short-sighted!" You may criticize me. You may believe that even though your results cannot lead to practical breakthroughs directly, they belong to the ongoing accumulation of knowledge about the brain. As the accumulation continues, people will eventually have a very good understanding of the brain and develop advanced brain technology to solve practical problems.

However, your criticism neglects a possibility: Some aspects of the brain may be unknowable. If such is the case, we may never be able to fully understand the brain, regardless of how much knowledge we accumulate. This unknowability reflects the fundamental limitations of human reasoning capabilities.

If you ever doubt the limitations of human reasoning, take a trip to an art museum. As you peruse the galleries, you may ask yourself: Is it possible to develop a logical system that, through a series of if-then reasoning, could lead to the creation of a masterpiece? If you doubt the existence of such a logical system for the creation of art, then why do you believe that a logical system for the workings of the brain exists? After all, the brain is believed to be much more complex than any artwork created by humans.

In this paper, I will review literature that highlights the limitations of mechanism-investigating research in solving practical problems. I will then explore the concept of the unknowability of the brain through the lenses of neuroscience, philosophy, physics, and AI. Finally, I will provide suggestions for conducting meaningful research in light of this unknowable reality.

Introduction

Atomism, the idea that the universe is composed of fundamental components known as atoms, is perhaps the most influential philosophy leading scientific research. Richard Feynman considered atomism to be the most important thinking we should pass on to the next generation (Feynman et al., 2011), as various physical changes and chemical reactions can be explained by supposing the movements and interactions of atoms (Feynman et al., 2011).

Atomism has also had a strong influence on cognitive neuroscience. Psychologists have divided cognition into several elements, including perception, learning, memory, and decision-making (Baldwin, 1893). Each of these elements can be further divided into several sub-elements from different perspectives. For example, perception can be divided into the perception of space and time or into visual and auditory perception. Memory can be divided into short-term and long-term memory, or episodic and semantic memory, among other things. After investigating the brain activity when the subject is performing each element of cognition, neuroscientists aim to understand the biological backend of cognition by collecting all these pieces together (Fig. 1a). From this atomistic perspective, studying a single cognitive element is the foundation for understanding cognition, which is why I name this research protocol to be *basic*.

To perform basic cognitive studies, researchers elaborately designed simple and well-controlled experimental conditions to study a single cognitive element while teasing apart the influence from other elements. For example, to study working memory, researchers trained monkeys to recall a visual cue after a delay period (Constantinidis et al., 2001) (Fig. 1b). To study decision-making, researchers trained monkeys to watch two types of dots moving toward opposite directions and then decide which type had more dots (Roitman & Shadlen, 2002) (Fig. 1c).

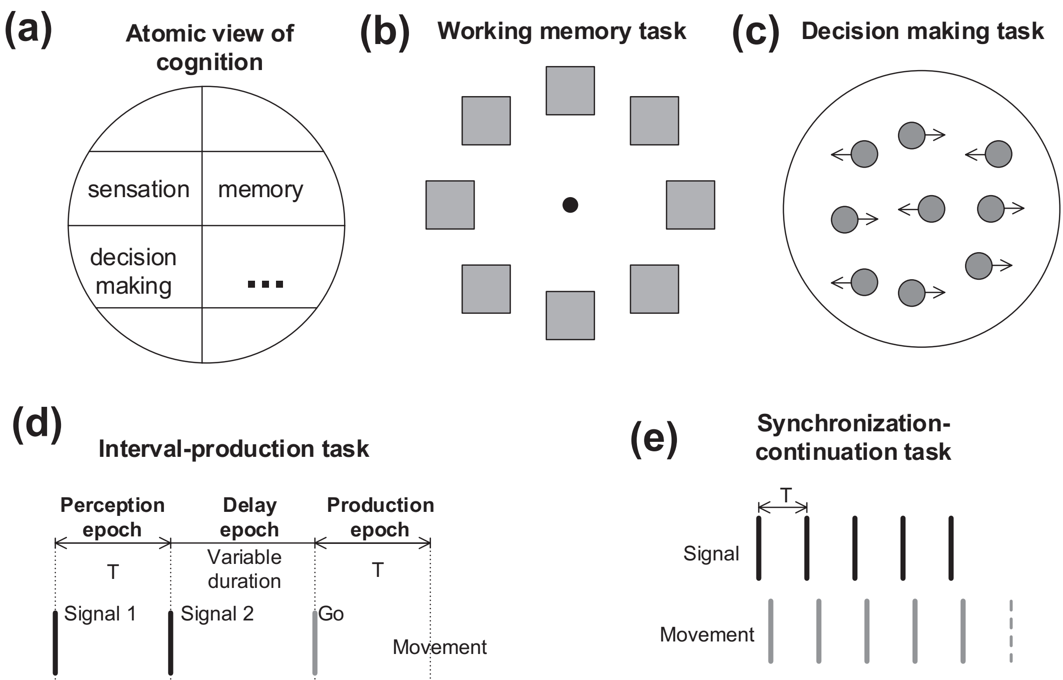

Fig. 1 Illustration of basic cognitive studies. (**a**) Basic cognitive studies are guided by the philosophy of atomism, which divides cognition into many elements, each of which is studied separately. Atomists believe that by understanding each element, we can eventually understand the whole of cognition. (**b**) A classical experiment to study working memory. The subject fixates on a central point, and a visuospatial cue (one of the eight gray boxes) is presented briefly, followed by a mnemonic delay. After the delay, the subject must make a saccadic eye movement to the remembered location. (**c**) A classical experiment to study decision-making. There are two types of random dots, one moving leftward and the other moving rightward. The subject must decide which type has more dots. (**d**) Schematic of the time production task. The subject receives two signals (black bars) separated by a time interval T; after a delay epoch with variable duration, a go cue (gray bar) appears, and the subject must move at time T after the go cue. (**e**) Schematic of the synchronization-continuation task. The subject must move (gray bars) immediately following a sequence of signals (black bars) with period T. The subject must still move with period T (dashed bar) after the signal was removed

Another example is the study of time cognition, which also stems from atomism. To focus on the processing of time while disentangling other cognitive elements (such as the perception of spatial information), psychologists or neuroscientists train subjects to perform simple but deliberately designed timing tasks. In a classical experiment (Rakitin et al., 1998), participants were presented with specific time intervals delimited by stimuli and then were asked to reproduce the interval (Fig. 1d). When subjects were performing these simple and deliberately designed tasks, researchers recorded subjects' brain activity to propose neural network mechanisms underpinning basic elements of cognition.

While some scientists think that the pure aim of science is to satisfy our curiosity about the world, I believe that scientific results must be implemented in technology and benefit the mass of people before scientific results complete their mission. However, the status and prospects of the technological applications of basic cognitive studies have seldom been discussed. In this paper, I will discuss the technological applications of basic cognitive studies, starting with a review of cognitive studies of time processing in the brain (i.e., basic timing studies) as examples of basic cognitive studies. Time processing is an indispensable dimension of cognition (Merchant et al., 2013), and time is also an indispensable dimension of any real-world signal to be processed by technology. Therefore, one may expect that the results of basic timing studies lay down the foundations for processing temporal signals in brain-related technology. Unfortunately, after reviewing two fields of brain-related technology, neuroengineering for brain health and brain-inspired artificial intelligence, which are two promising application fields of neuroscience suggested by the China Brain Project (Poo et al., 2016), I found that the results of basic timing studies are hardly helpful in solving practical problems.

I will attempt to clarify this awkward situation and offer suggestions for future research (Fig. 2). In my view, the challenge of applying basic timing studies (and, more broadly, basic cognitive studies) to technology stems from their lack of generalizability. In other words, the results of these studies are contingent on the specific conditions and tasks of the laboratory experiments that produced them and may not be applicable in other contexts. This lack of generalizability may be rooted in the fundamental unknowability of the world, including cognition. In other words, the capability of knowledge to describe the world is fundamentally limited, so the generalizability of our knowledge to various situations in the world is fundamentally limited, and therefore, the capability of knowledge to guide technological creation to change the world is also fundamentally limited.

I suggest three ways to improve future research (Fig. 2). Firstly, to improve the generalizability of results, researchers should analyze brain activity in real-life settings, rather than simple tasks in well-controlled experimental conditions, and examine their results under various situations. Secondly, to deal with the unknowability of the world, researchers should engineer surrogates of the object under investigation, so that they can predict the behavior of the investigated object using the surrogate, even without understanding how the object under investigation works. Finally,

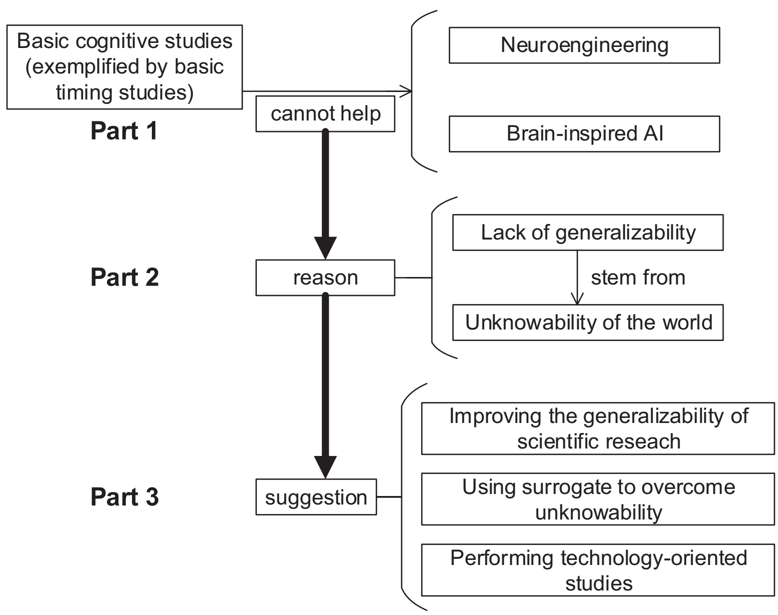

Fig. 2 Overview of this paper. In Part 1, the main results of basic timing studies and two application fields (neuroengineering and brain-inspired AI) are reviewed, showing that basic timing studies (and perhaps more generally, basic cognitive studies) cannot help the application fields of neuroscience. In Part 2, it is proposed that this situation is due to the lack of generalizability of the basic timing studies and, more fundamentally, the unknowability of the world. Finally, in Part 3, researchers are suggested to improve the generalizability of their results, engineer surrogates to overcome the unknowability of the world, and perform technology-oriented studies

due to the fundamental unknowability of the world, I suggest that researchers focus on technology-oriented research with the aim of creating new technologies, rather than science-oriented research with the aim of discovering knowledge.

Basic Timing Studies

This section provides an overview of basic timing studies, which use two main paradigms to study time cognition. The first paradigm is interval timing, which involves training the subject to perceive or produce a single time interval (Fig. 1d). The second paradigm is beat timing, which involves training the subject to perceive or produce a sequence of time intervals rhythmically (Fig. 1e). In the interval-production task (Rakitin et al., 1998), an example of the first paradigm, the subject is presented with a specific time interval delimited by stimuli and is then asked to reproduce the interval (Fig. 1d). In the synchronization-and-continuation task (Gámez et al., 2019), an example of the second paradigm, the subject is required to act following a sequence of rhythmic stimuli and continue to act rhythmically after the removal of the stimuli (Fig. 1e). By recording brain activity during these tasks, researchers can discover features of brain dynamics related to time cognition and gain insight into the neural network mechanisms underpinning time cognition.

In the interval-production task (Fig. 1d), a neural network perceives time by evolving its state along a stereotypical trajectory in the perception epoch, maintains time intervals in working memory using a manifold of line attractor in the delay epoch, and predicts a coming event by evolving its state along isomorphic trajectories with the speed of state evolution inversely scaling with the to-be-produced time interval in the production epoch (Bi & Zhou, 2020a) (Fig. 3a). These dynamic features align with experimental findings from other interval-timing tasks (Jin et al., 2009; Mita et al., 2009; Wang et al., 2018). In the synchronization-continuation task (Fig. 1e), the network encodes different beating periods T using circular trajectories (Gámez et al., 2019) (Fig. 3b). The radii of these circular trajectories increase with the period T, but the speed of state evolution with time remains constant across different values of T.

In both the perception and production epochs of the interval-production task, as well as in the beating intervals in the synchronization-continuation task, the neural network relies on state evolution along trajectories to sense the passage of time. This state evolution can be achieved through several mechanisms, including the pacemaker-accumulator model (Buhusi & Meck, 2005) (recently supported in (Cook et al., 2022)), in which an accumulator counts the number of pulses received from a pacemaker (Fig. 3c, left); the synfire chain model (Zeki & Balci, 2019), in which a chain of neurons is sequentially excited (Fig. 3c, middle); and the striatal beat-frequency model (Matell & Meck, 2004), in which a group of oscillators with heterogeneous frequencies have their phases reset by the stimulus (Fig. 3c, right).

Anatomically, several brain areas have been identified as participating in timing, including the basal ganglia (Jin et al., 2009), supplementary motor area (SMA) (Mita et al., 2009), sensory cortex (Shuler & Bear, 2006), and prefrontal cortex (Wang et al., 2018). There is ongoing debate about whether timing relies on dedicated circuits in the brain or on intrinsic computation that emerges from the inherent dynamics of neural circuits (Paton & Buonomano, 2018; Ivry & Schlerf, 2008). A prevailing viewpoint is that timing depends on the interaction of core timing areas, such as the basal ganglia and SMA, which are consistently involved in temporal processing across various contexts, and other areas, such as the prefrontal cortex, sensory cortex, and cerebellum, which are activated in a context-dependent manner (Merchant et al., 2013).

At the behavioral level, the most well-known timing principle is the scaling property, which posits that the variance of time interval estimation is proportional to the mean of the estimation (Allman et al., 2014).

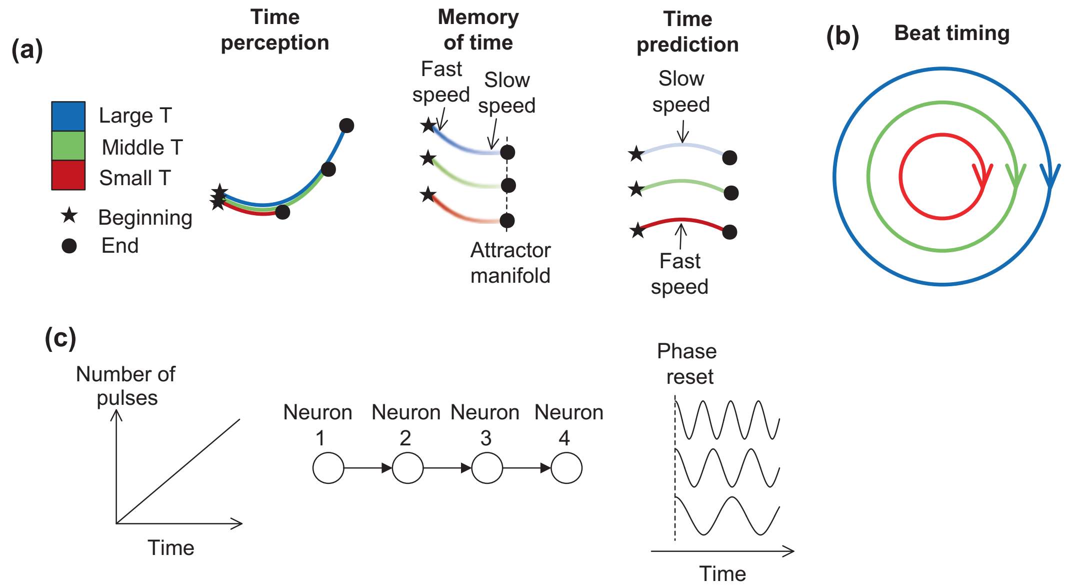

Fig. 3 Some results of basic timing studies. (**a**) Dynamic features of neural networks in the interval-production task. Left: For time perception in the perception epoch (see Fig. 1d), the network exhibits a stereotypical trajectory whose final position determines the perceived time interval T (see Fig. 1d). Lines with blue, green, and red colors, respectively, represent the trajectories when T is large, middle, and small. Asterisk and circle, respectively, represent the beginning and end of the trajectory. Middle: In the delay epoch, time intervals are maintained in the working memory as positions (black dots) in an attractor manifold. The speed of state evolution with time decreases near the attractor (indicated by the lighter color near the attractor). Right: In the production epoch, time prediction is performed when the network state evolves along isomorphic trajectories, with the speed of state evolution inversely scaling with the to-be-produced interval T. (**b**) Dynamic features of neural networks in the synchronization-continuation task. With different periods T (see Fig. 1e), the network state evolves along different circular trajectories at the same speed, but the radius of the circular trajectory increases with T. (**c**) Some computational models for the neural mechanisms of time sensing. Left: In the pacemaker-accumulator model, time is measured by the accumulated number of pulses emitted from the pacemaker. Middle: In the synfire chain model, time can also be measured by the sequential firing of a chain of neurons. Right: In the beat-frequency model, time is measured by the activity pattern of a group of oscillators with heterogeneous frequencies after phase resetting

Brain-Related Technology

Time processing is a fundamental aspect of cognition (Merchant et al., 2013), and time is also an indispensable dimension of any real-world signal to be processed in technology. Therefore, one might expect that studies on time processing in cognition would profoundly influence brain-related technology. This section will review two fields of brain-related technology, neuroengineering for brain health and brain-inspired artificial intelligence, which are two promising application fields of neuroscience suggested by the China Brain Project (Poo et al., 2016). Unfortunately, we will see that the results from basic timing studies are hardly helpful in solving practical problems.

Neuroengineering for Brain Health

Neuroengineering involves designing interfaces between living neural tissue and non-living constructs in order to understand, repair, replace, or enhance neural systems (Hetling, 2008). In this paper, I will review neuroengineering techniques used for the therapy of Parkinson's disease through deep brain stimulation, the diagnosis of epilepsy through neuroimaging, and the development of speech prostheses through machine translation of brain activity into language.

Parkinson's disease is closely related to pathological changes in the basal ganglia (Poewe et al., 2017), a core timing area of the brain (Merchant et al., 2013). Epilepsy also recruits timing-related regions such as the thalamus, basal ganglia, and frontal lobe (Bertram, 2009; Wu et al., 2019). As a result, patients with either Parkinson's disease or epilepsy may experience distortion of timing perception (Gu et al., 2016; Greyson et al., 2014; Cainelli et al., 2019). Besides, language has rich hierarchical temporal structures, and the processing of language may share a similar neural substrate with the processing of music (Patel, 2003; Janata & Grafton, 2003; Hickok, 2012). Therefore, it is reasonable to assume that basic timing studies could be of great help in the therapy and diagnosis of Parkinson's disease and epilepsy, as well as the machine translation of brain activity into language. However, I will show below that this is not the case.

Neuroengineering Is Driven by Clinical Data and Experience

Deep brain stimulation (DBS) therapy for Parkinson's disease was pioneered by Lawrence Pool, who implanted an electrode into the caudate nucleus of a female patient in 1948 (Pool, 1954). Traditional DBS is an open loop, where the clinician sets parameters of the controller that deliver short-duration (60–180 ms) and high-frequency (typically 130–185 Hz) pulses of electrical stimulation to alleviate symptoms (Benabid et al., 1994; Limousin et al., 1995; Siegfried & Lippitz, 1994) (Fig. 4a, left). However, this type of DBS cannot adapt its stimulation according to the feedback from the patients and has several drawbacks, such as adverse effects such as dyskinesia and high battery consumption (Bouthour et al., 2019; Krauss et al., 2021). Recently developed closed-loop DBS overcomes these problems by delivering stimulation only when pathological biomarkers are detected (Bouthour et al., 2019; Krauss et al., 2021) (Fig. 4a, right).

Interestingly, despite the broad success and application of DBS, the mechanism by which DBS ameliorates Parkinson's disease is still not fully understood, although some mechanisms related to neuronal circuits, astrocytes, and neurogenesis have been proposed (Okun, 2012). Due to the lack of understanding of the mechanism, the technical details of DBS have been established mainly through empirical means. For example, the optimal stimulation waveform shape in open-loop DBS was determined by systematically varying stimulation parameters and examining the therapeutic effects (Rizzone et al., 2001; Kuncel et al., 2006). The most prominent biomarker used in closed-loop DBS, excessively synchronized beta oscillation, was also discovered through empirical comparisons between normal and diseased brains (Oswal et al., 2013; Cheyne, 2013). Therefore, mechanical insight, which is the aim of basic timing studies (Fig. 3), is not the primary driving force behind the development of DBS.

Though the mechanical insights provided by basic timing studies may not currently be helpful in the research of DBS, one might still expect that they could be useful in the future. However, recent research trends suggest a dominance of data-driven automatic design in the development of DBS technology, rather than a rational implementation of mechanical knowledge. As mentioned earlier, closed-loop DBS delivers stimulation into the brain only when pathological activities (i.e., biomarkers) are detected. Traditionally, excessive beta oscillation was predetermined as the key biomarker of patients' tremors in Parkinson's disease (Bouthour et al., 2019; Krauss et al., 2021). However, in two recent studies (Shah et al., 2018; Tan et al., 2019), the authors recorded patients' body movements using accelerometers and recorded local field potentials (LFPs) using electrodes. They then trained binary classifiers to detect the LFPs during tremor or non-tremor periods. Here, the detector (i.e., the binary classifier) is trained by clinical data, instead of being rationally designed using our knowledge of the mechanisms of Parkinson's disease. A similar data-driven approach has also been used to detect biomarkers of depression (Scangos et al., 2021a, b), where a classifier was trained to map stereoelectroencephalography (SEEG) recordings to depression scores measured by a psychological questionnaire.

This data-driven approach is also the mainstream of other neuroengineering techniques. For example, in recent studies on epilepsy diagnosis, neural network models were built to simulate the large-scale dynamics of the brain. The models

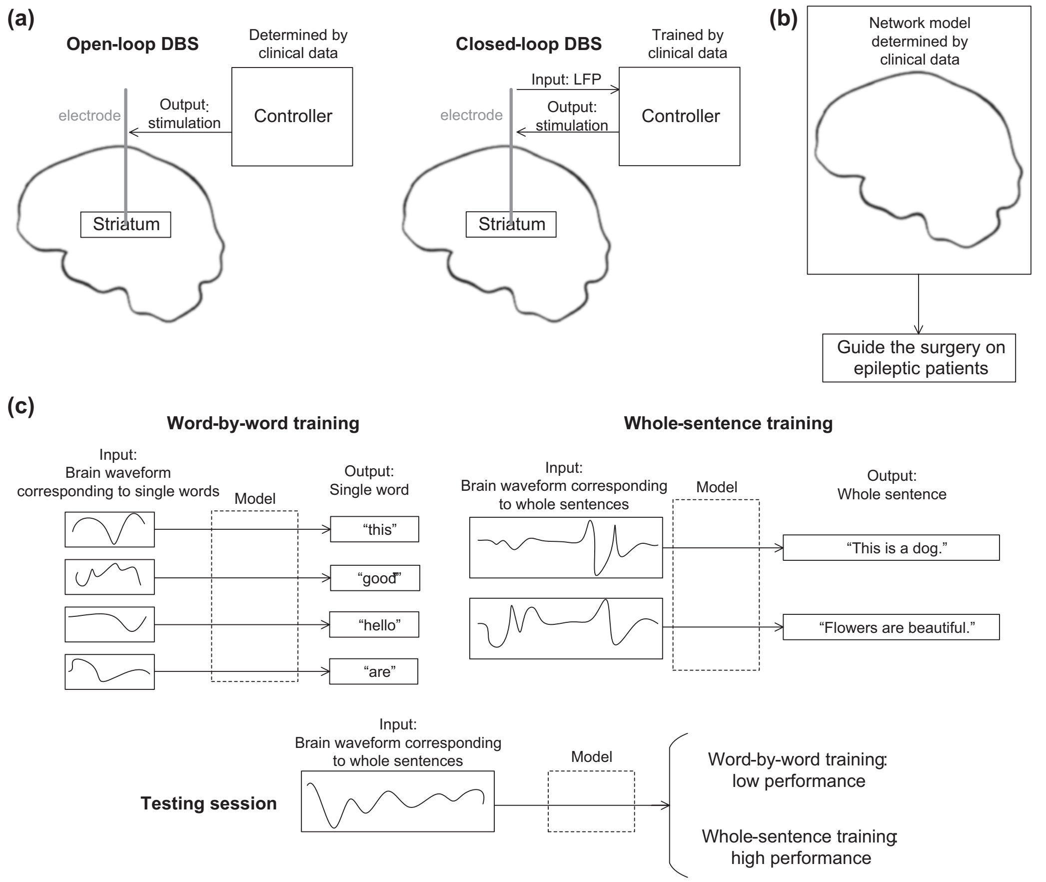

Fig. 4 Some neuroengineering techniques. (**a**) In open-loop deep brain stimulation (DBS) (left), the stimulation waveform is preset in the controller. In closed-loop DBS (right), the stimulation waveform can be adjusted according to the local field potential (LFP) of the brain. In both types of DBS, the controller is programmed based on clinical data, rather than on mechanical understandings from basic cognitive studies. (**b**) Neural network models that simulate large-scale dynamics of the brain have been used to guide the surgery on epileptic patients. The parameters of the model are also determined by clinical data, rather than mechanical understandings from basic cognitive studies. (**c**) There are two strategies to build the model to translate brain activities to language for speech prosthesis. In the word-by-word training strategy (upper left), the model is trained to map the brain waveforms when the patient is speaking single words to single spoken words. In the whole-sentence strategy (upper right), the model is trained to map the brain waveforms when the patient is speaking whole sentences to the whole spoken sentences. In the test session (lower), the trained model is used to translate the brain waveforms corresponding to whole sentences. The whole-sentence training strategy results in better performance than the word-by-word strategy

were used to identify the ictogenic zone of seizures and guide the resection of brain areas in clinical surgery (Cao et al., 2022; Sinha et al., 2017) (Fig. 4b). In their neural network models, the connection strengths were determined through the fitting of empirical data, rather than being rationally designed based on mechanical insights into epilepsy or the information processing of the brain. Another example is the machine translation of brain activities to language, which can be used as a speech prosthesis for degenerative motor diseases such as amyotrophic lateral sclerosis and locked-in syndrome (Fig. 4c). Traditional approaches trained translation machines by mapping neuroimaging signals to individual words or even sub-word syllabic features (such as vowel harmonics and fricative consonants) (Pasley et al., 2012; Angrick et al., 2019) (Fig. 4c, upper left). However, the best brain-to-language translation performance is now realized by training recurrent neural networks in an end-to-end manner, mapping brain signals to entire sentences rather than single words or syllabic features (Makin et al., 2020; Cogan, 2020; Moses et al., 2019) (Fig. 4c, upper right). Although understanding the brain activities related to words and even syllables may seem more "basic" and "mechanical," implementing such understandings in the technique results in worse performance than directly training a neural network to map brain signals to entire sentences (Fig. 4c, lower).

Summary

Overall, basic timing studies (Fig. 3), though expected to be "basic," unfortunately do not lead to the progress of application-oriented research. This awkward situation is due to a gap in that basic timing studies aim for mechanistic explanations for simple timing tasks, but neuroengineering, which aims for good performance in practical use, is not mainly driven by mechanistic understandings of the brain, but by clinical data and experience. This gap not only exists between basic timing studies and neuroengineering, as we have discussed here, but, more generally, between basic cognitive studies (Fig. 1) and neuroengineering. Therefore, we may conclude that basic cognitive studies do not lead to the progress of neuroengineering.

Brain-Inspired Artificial Intelligence

Brain-inspired artificial intelligence (AI) is another potential application field of neuroscience. Brain-inspired AI aims to build strong AI (i.e., AI that has mental capabilities and functions that mimic the human brain, or in other words, can pass the Turing test) by mimicking the structure and function of the brain, through the implementation of neuroscience knowledge in AI engineering (Hassabis et al., 2017). I have a criticism of this brain-inspired approach to AI, though due to ethical concerns, detailed investigations on the human brain cannot be performed. As a result, brain-inspired AI can only closely mimic the brain of animals, which has low-level intelligence, rather than that of humans, whose high-level intelligence is the ultimate aim. Therefore, the brain-inspired approach should not be the leading approach to strong AI in the long run. I will talk about the possible approach to strong AI at the end of this subsection; at present, however, let us forget this criticism and think about how basic timing studies may contribute to brain-inspired AI. Unfortunately, I will show that basic timing studies are also of little help to this field.

The Inspiration for AI from Neuroscience

The inspiration for AI from neuroscience is found at the levels of neurons, synapses, and neural networks. This is exemplified below:

1. Single Neuron Level.
   (a) Biological neurons fire spikes, unlike artificial analog neurons, whose activities take continuous values. Implementing spiking neurons in hardware significantly reduces energy consumption compared to analog neurons (Frenkel, 2021). The reason is that the membrane voltage of spiking neurons stays near the resting state most of the time due to the sparsity of spiking periods, resulting in small leaky currents.
   (b) Biological neurons also have rich internal dynamics due to the interaction between the membrane voltage and ion channels (Dayan & Abbott, 2001), unlike artificial neurons, which are usually nonlinear filters of total synaptic currents. Such rich internal dynamics significantly improve the computational power of biological neurons (Beniaguev et al., 2021). Recently, it has been found that only 19 neurons with internal dynamics can make up a full-stack autonomous vehicle control system (Lechner et al., 2020).
2. Single Synapse Level.
   (a) Biological synapses have binary efficacies (O'Connor et al., 2005), unlike in artificial networks where synaptic weights typically take continuous values. Binary-weight artificial neural networks have

been investigated and broadly used due to their low computation and memory cost, as well as performance that is comparable with continuous-weight networks (Courbariaux et al., 2016).

(b) Biological synapses also have hidden states other than synaptic efficacy, which arise from the complex interactions of proteins in synapses (Graupner & Brunel, 2010). Adding hidden synaptic states in artificial neural networks improves memory capacity and learning performance (Baldassi et al., 2007; Kirkpatrick et al., 2017). The reason is that a high hidden state of a synapse can indicate that this synapse is important for the good performance of a task; therefore, protecting the efficacy of synapses with high hidden states from being changed in the further training process can maintain the performance of the neural network during further training.

3. Neural Network Level.

(a) Memory replay, found in the hippocampus and cortex (Ji & Wilson, 2007), is a phenomenon in which the neuronal firing sequence in sleep or at rest closely matches the firing sequence in the real experience just before. Memory replay inspires DQN (Mnih et al., 2015), a well-known deep reinforcement learning algorithm that guides actions according to perceptual inputs in order to maximize future rewards. Besides, memory replay is also used in the Dyna algorithm (Sutton & Barto, 2018) to train a mental model of the environment. After training, the agent can predict the outcome of an action in situations never seen before using this mental model, facilitating the agent to adapt to more complicated environments.

(b) Biological neurons are subject to gain modulation, which means that one input, the modulatory one, affects the sensitivity of a neuron to another input (Salinas & Thier, 2000; Salinas & Sejnowski, 2001). Gain modulation is the neural mechanism of attention. With attention mechanisms, a neural network looks at an image or input sequence and decides which parts of the image or sequence are important for the task at hand and then sends only the important parts to subsequent information processing. Attention mechanisms have become an indispensable component of advanced image and language processing models (Vaswani et al., 2017; Devlin et al., 2019).

(c) Context-dependent gating (Cichon & Gan, 2015) means that different sparse sets of dendritic branches are disinhibited when the brain is involved in different tasks. This mechanism allows the brain to recruit different dendritic branches for different tasks, so that the synaptic weights learned for one task will not interfere with the configuration learned for another task. Such context-dependent gating has been implemented in artificial neural networks to avoid catastrophic forgetting during continual learning (Manning et al., 2020; Zeng et al., 2019).

Basic Timing Studies Hardly Inspire AI

From the examples provided (also see (Hassabis et al., 2017) for a detailed review), it is clear that basic timing studies do not have a significant impact on the development of brain-inspired AI, despite time processing being a fundamental aspect of brain cognition. Similarly, other basic cognitive topics, such as working memory and decision-making (Fig. 1b, c), though attract great interest in the neuroscience community, and they also contribute little to brain-inspired AI. There are two possible reasons for the limited impact of basic cognitive studies in AI applications:

1. Lack of Generalizability (Fig. 5a).

   All the neural mechanisms implemented in AI have a common property: They are not task-specific. In other words, if a neural mechanism exists only when the brain is performing a simple task like Fig. 1b–e, but does not exist if the brain is performing another more complicated task, this neural mechanism will not be used in AI implementation. The reason

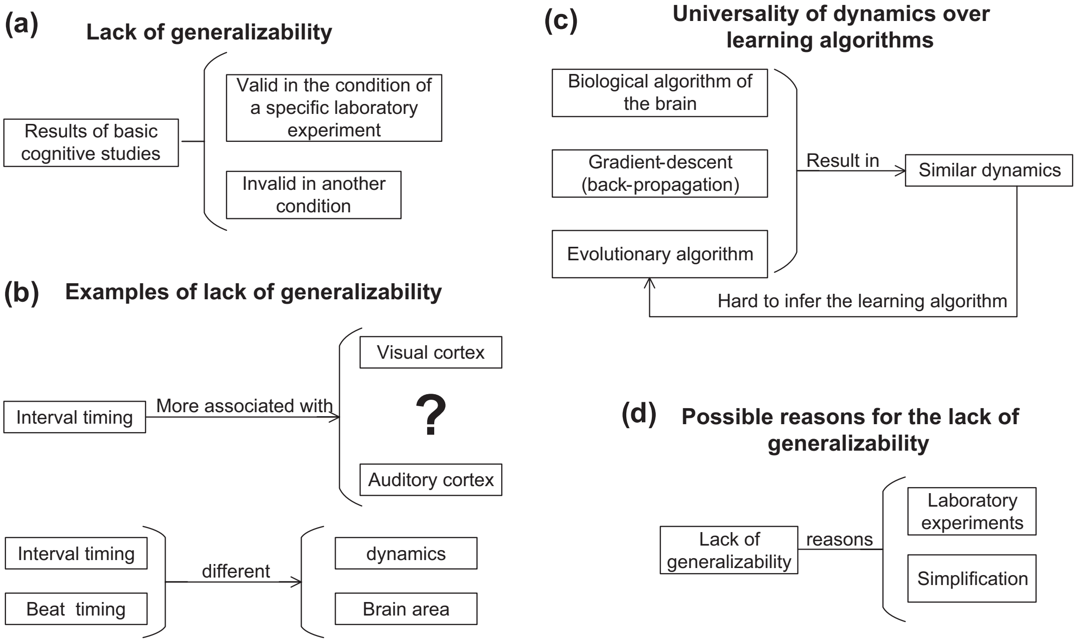

Fig. 5 Pitfalls of basic cognitive studies. (**a**) Basic cognitive studies may lack generalizability, so that the results are valid in a specific laboratory experiment condition but invalid in another condition. (**b**) Examples of lack of generalizability. Upper: Some experiments suggest that interval timing is more associated with the visual cortex, but other experiments suggest that interval timing is more associated with the auditory cortex. Lower: Interval timing and beating timing, though both are pure timing tasks, have different dynamic features and involve different brain areas. (**c**) Different learning algorithms result in similar neuronal population dynamics when the neural network is trained on the same basic cognitive task. Therefore, we cannot infer the learning algorithm that the brain uses through the dynamics when the brain is performing basic cognitive tasks. (**d**) The lack of generalizability may be due to the methodology of laboratory experiments in basic cognitive studies and also the misunderstanding of the ideology of simplification

is simple: AI aims to solve complicated real-world problems, instead of toy problems like Fig. 1b–e designed by neuroscientists.

The lack of generalizability to real-world situations is the main shortcoming of basic cognitive studies. The dynamics of the brain when performing complex tasks cannot be deduced from the dynamics observed when the brain is performing simple tasks. In other words, even if we have a good understanding of the dynamics involved in numerous simple tasks like the one illustrated in Fig. 1b–e, we still do not know the brain dynamics in complex tasks. For example, suppose we let a patient perform simple tasks of speaking single words. Even if we record the brain activity related to numerous single-word speaking, we still do not know the patient's brain activity when speaking a whole sentence: because by dividing a sentence into single words, we are neglecting the syntactic structure of the sentence. This is why the translation of brain activity to language for speech prosthesis achieves better performance when the neural network is trained to translate one sentence at a time instead of one word at a time (Makin et al., 2020; Cogan, 2020; Moses et al., 2019) (Fig. 4c). More generally, cognition requires the coordination of all "basic" elements: perception, memory, decision-making, and so on. Even if we study each "basic" element in isolation, we will still not be able to understand how the brain performs complex real-world tasks that require the coordination of these elements. We will discuss more about the lack of generalizability further in the next section.

2. Lack of Insight into Brain Learning Mechanism (Fig. 5c).

   One may wonder whether the dynamic features observed when the brain performs simple tasks (Fig. 3) depend on the specific learning algorithm of the brain. If a particular learning algorithm results in the experimentally observed features (Fig. 3), while other algorithms do not, we may be able to infer the brain's learning algorithm through these dynamic features. This learning mechanism could then be implemented into AI design. Unfortunately, accumulating evidence suggests that similar dynamic features universally emerge when neural networks are trained on the same basic cognitive task using different learning algorithms. This implies that we may not be able to infer critical information about the learning algorithm that the brain uses through the dynamics observed when the brain is performing simple tasks.

   For example, although error back-propagation (BP) algorithms lack convincing experimental support (Lillicrap et al., 2020), artificial neural networks trained by BP exhibit biologically plausible dynamics in image classification tasks (Hong et al., 2016), language tasks like next-word prediction (Goldstein et al., 2022), and other simple tasks in basic cognitive studies (Bi & Zhou, 2020a; Mante et al., 2013). Recently, I trained recurrent neural networks using an evolutionary algorithm to perform the context-dependent decision-making task (Bi et al., 2022) and found that the resulting network exhibited dynamics closely analogous to those observed in monkey experiments and the dynamics observed in artificial neural networks trained by BP (Mante et al., 2013). The reason for this universality of dynamics across different learning algorithms is unknown, but it is possibly because different algorithms universally tune the synaptic weights into a high-entropy region in the synaptic configuration space (Baldassi et al., 2015; Bi & Zhou, 2020b). Here, "high entropy" means that if we slightly perturb the synaptic weights found by an algorithm, the perturbed weights still probably result in good task performance. Therefore, the weights found by different algorithms are likely to be close to each other in a high-entropy region, which may be the reason for the universal dynamic property of the networks trained by different algorithms. Due to this universality, we cannot gain insight into the learning mechanism of the brain from the dynamic features found in basic cognitive studies, let alone implement the brain learning mechanism in AI.

So what is the approach to strong AI, the machine with intelligence equal to the human brain or even more powerful? In my opinion, the most important thing we should learn from biology is the colossal scale of the human brain. Comparative studies have shown that the human brain contains more neurons than any other animal, which is probably the reason for our superior cognitive abilities (Herculano-Houzel, 2012). Consistently, AI is undergoing a paradigm shift with the rise of colossal models (e.g., BERT (Devlin et al., 2019) and GPT-3 (Brown et al., 2020)) with over 100 billion parameters trained on oceans of data (Bommasani et al., 2021). Such models, trained unsupervisedly, develop geometric representation of knowledge (Manning et al., 2020; Rives et al., 2021), which versatilely serve as the common basis of many task-specific models via adaptation (Bommasani et al., 2021). Most impressively, as the size of the neural network increases, advanced functionalities such as in-context learning naturally emerge (Brown et al., 2020): In-context learning means that the neural network after training can be competent for a task never seen during training, after the trained neural network is instructed by a natural language description of the task. Such colossal models are becoming the trend of AI led by big tech companies such as Google, Microsoft, and Huawei, with broad applications in the text (Devlin et al., 2019), images (Ramesh et al., 2021), protein design (Rives et al., 2021), and chemical reactions (Schwaller et al., 2021). Recently, at a conference, a manager of the colossal model project of Huawei told me that the progress of colossal models is also gradually getting stagnant, because we cannot afford the huge energy consumption to train the model if the model is too large (Strubell et al., 2019). He believed that the next generation of colossal models should be led by the revolution of the computation paradigm, such as quantum computing, which can speed up some kinds of computations by exponential order (Nielsen & Chuang, 2011). At another conference, a professor of quantum physics told me that quantum algorithms are getting mature; the bottleneck of quantum computing lies in its hardware implementation. If their opinions are correct, we may expect that strong AI will naturally emerge after the manufacturing technology of quantum computers is mature and if we train mega-colossal models using quantum computers.

Summary
Overall, basic cognitive studies cannot significantly contribute to AI due to their limitations on generalizability and insight into brain learning mechanism. The future of AI is likely to be led by colossal models.

Contemplation

Time processing is a fundamental aspect of cognition, and time is also an indispensable dimension of any real-world signal to be processed in technology. But why do basic timing studies, which aim to study time processing in the brain, which are also interesting and elegant, instead have little help to the progress of brain-related technology? Below, I will discuss possible reasons for this discrepancy.

Generalizability: The Shortcoming

Generalizability is a measure of how useful the results of a study are for broader situations. Generalizability is the critical hypothesis (and also the aim) of science. To understand this point, let us consider a simple example. If we want to test the effectiveness of a new drug, we will recruit several patients and test the drug on them. However, our aim is not only to investigate these several recruited patients, but to draw a general conclusion on the effect of the drug on the mass of people using these recruited patients, with the hypothesis that similar phenomena can also be observed if we recruit another group of patients. As another example, when physicists perform an experiment and conclude a physical law, their aim is not only to explain the very experiment they perform, but to conclude a law generalizably applicable to other experiments taken at another place and another time. However, we should not take such generalizability for granted. Many hard problems are because we do not have a generalizable understanding of the problem or a generalizable technique to deal with the problem. For example, cancer is a challenging disease to cure because we do not have a generalizable technique to efficiently kill all the cancer cells due to the high diversity of cancer cells (Morita et al., 2020; Black & McGranahan, 2021).

Lack of generalizability is a significant shortcoming of basic timing studies. Results obtained under one experimental condition often cannot predict the result under another condition. For example, if an auditory stimulus is associated with time duration T_a and a visual stimulus with duration T_v in a rat subject, presenting the auditory and visual stimuli simultaneously will make the rat subject time an expected duration of T_+, which is between T_a and T_v, but closer to T_v (Swanton & Matell, 2011; Matell & Kurti, 2014). Additionally, compared to an auditory stimulus, the association between a visual stimulus with a time duration can be better transferred to a subsequent operant response when tested in a Pavlovian-instrumental transfer procedure (Matell & Valle, 2017). These results imply that visual signals are more involved in interval timing than auditory signals. However, in a recent study on an action timing task, in which a mouse had to learn the timing of its action based on the sensory feedback caused by its own action, it was the deprivation of auditory input (not visual) that disrupted the learned action timing (Cook et al., 2022), contradicting previous understanding (Fig. 5b, upper). Furthermore, there are two frequently studied experimental paradigms of timing tasks: interval timing, in which the subject is to perceive or produce a single time interval (Wang et al., 2018; Karmarkar & Buonomano, 2007) (Fig. 1d), and beating timing, in which the subject is to perceive or produce regular beats (Gámez et al., 2019) (Fig. 1e). It has been found that the brain uses different neural substrates and mechanisms to process temporal information in these two paradigms (Gámez et al., 2019; Wang et al., 2018; Karmarkar & Buonomano, 2007; Teki et al., 2011), even though they are both pure timing tasks with no other information (such as spatial information) involved (Fig. 5b, lower). Overall, the lack of generalizability in basic timing studies makes it challenging to conclude how the brain processes temporal information.Below, I will explore two possible reasons for the lack of generalizability in basic timing studies (Fig. 5d):

1. Laboratory Experiments.

   Basic timing studies are performed in laboratory experiments, which have artificially designed and well-controlled experimental conditions (just like Fig. 1d, e) that may not reflect real-life situations. The lack of generalizability has long been recognized as the shortcoming of laboratory experiments in social science, including psychology (Brüggemann & Bizer, 2016; Hulstijn, 1997). Therefore, the limitations of basic timing studies discussed here are just examples of the general shortcoming of the laboratory experiment paradigm. Perhaps the only way to improve the generalizability of the results of laboratory experiments is to capture the common results of different experimental conditions through a literature review. For example, Bueti and Buonomano (2014) concluded that temporal learning transfers across different modalities, including visual and auditory modalities, different auditory pitches, and slightly different lengths of temporal intervals, by reviewing papers. However, literature review cannot always lead to a straightforward conclusion, especially if the brain has a complicated performance under different experimental conditions. For example, the transfer of learning may not exist under some conditions and may be strong or weak in other conditions.

2. Misunderstanding of Simplification.

   Simplification is a pervasive idea in the data analysis and computational models of basic timing studies (Fig. 3). The pacemaker-accumulator model (Buhusi & Meck, 2005) (Fig. 3c, left), the best-known timing model, contains only four components (pacemaker, accumulator, memory device, and comparator) to model the timing process. The dynamic features found by basic timing studies (Fig. 3a, b) are often discovered after reducing the dimension of the population dynamics of neural networks using principal component analysis (PCA). This PCA method also manifests the idea of simplification: simplifying the population dynamics by reducing its dimension.

   The idea of simplification, also named the principle of Occam's razor, tries to explain the world using as few entities as possible. However, the advantage of this principle of simplification must be understood before using it. One widely accepted advantage of simplification is that simple theories tend to be more testable and, therefore, easier to falsify (Baker, 2016; Sober & Knowles, 1991). In other words, the primary advantage of a simple theory is not that it can better predict the experiment, but instead lies in its ease of falsification, which is believed to be a necessary property of a scientific theory (Popper, 1959). Another advantage of simplification (with controversy) is that it improves induction: choosing a simple theory after numerous observations reduces the chance of changing the theory after more future observations (Baker, 2016). This induction advantage is closely related to the concept of generalizability we discuss here because reducing the change in theory after future observations means improving the generalizability of the theory. However, "induction" means that the theory must be concluded after numerous observations, which is apparently not the case for the results (Fig. 3) in basic timing studies, which are usually proposed based on single laboratory experiments under simple and well-controlled situations. In other words, if we indeed want a simple timing theory that is generalizable to real-world situations, numerous observations in real-world situations are necessary.

Unknowability: The Reality

The methodology of basic cognitive studies involves recording brain activity while subjects perform deliberately designed tasks in order to understand the neural mechanisms of cognition. This methodology is based on the following philosophical understanding of science (Fig. 6a): Science investigates the world, generates knowledge, and then technology uses the knowledge generated by science to change the world. However, this philosophy fails to consider the possibility that the capability of knowledge (and therefore science) to describe the world may be fundamentally limited, meaning that some parts of the world are unknowable. If this is the case, the knowledge generated by science will not be able to well guide the design of technology to change the world effectively (Fig. 6a). The lack of generalizability discussed before may also stem from the unknowability of the world (including cognition): If the capability of knowledge to describe the world is fundamentally limited, we should not dream of the luxury that our knowledge has the possibility to generalize to every situation.

The Inspiration from AI

To discuss the limited capability of knowledge, let us start with an interesting empirical finding in AI. In AI, knowledge is usually represented by *(subject; relation; object)* triplets, representing the relationship between a subject and an object (Hogan et al., 2021). For example, the sentence "dog is animal" can be represented by a triplet *(dog; be; animal)*. A collection of a large number of triplets is called a knowledge base. It has been found that adding knowledge bases to deep learning models can improve the performance of natural language processing (Guo et al., 2022; Annervaz et al., 2018). However, interestingly, well-known colossal models (such as GPT-3 (Brown et al., 2020) of Microsoft or Pangu (Zeng et al., 2021) of Huawei) are pure deep neural networks without a knowledge base. A possible explanation for why well-known colossal models do not contain a knowledge base is that the performance improvement after adding a knowledge base to colossal models is marginal (below 4%, see Table 5 of Colon-Hernandez et al., 2021) (Fig. 6b). I discussed this interesting phenomenon with an AI expert in NetEase, who believed that this is because colossal models are trained by oceans of texts collected from the Internet, which contain far richer information than knowledge bases can provide, so adding knowledge bases to colossal models can hardly increase the information used to train the colossal models. Notice that people have invested great efforts to develop knowledge bases: Well-known knowledge bases such as YAGO and Freebase contain over 1 billion triplets. Despite such efforts, these knowledge bases are still hardly useful in the core AI technology of colossal models.

What can we learn from this empirical finding in AI? Notice that science is a process of generating knowledge from experiments (Fig. 6a): For example, basic timing studies aim to establish the relationship between the dynamics of the brain and the behavioral task. Also, notice that AI represents the future of technology. Therefore, if knowledge bases cannot help AI, we may conclude that science will not help technology in the future!

The Inspiration from Philosophy and Physics

The recognition of the limited capability of knowledge has a long history in philosophy. David Hume believed that causality cannot be justified because we can only observe that one thing, *A*, happened after another thing, *B*, but we cannot observe the underlying causal mechanism that made *A* happen after *B* (David Hume, https://en.wikipedia.org/wiki/David_Hume). Immanuel Kant believed that there exist things (the so-called things-in-themselves) that are unperceivable and unknowable. What we can perceive are mere "appearances" of these unknowable things, and a theory of the world develops when the per-

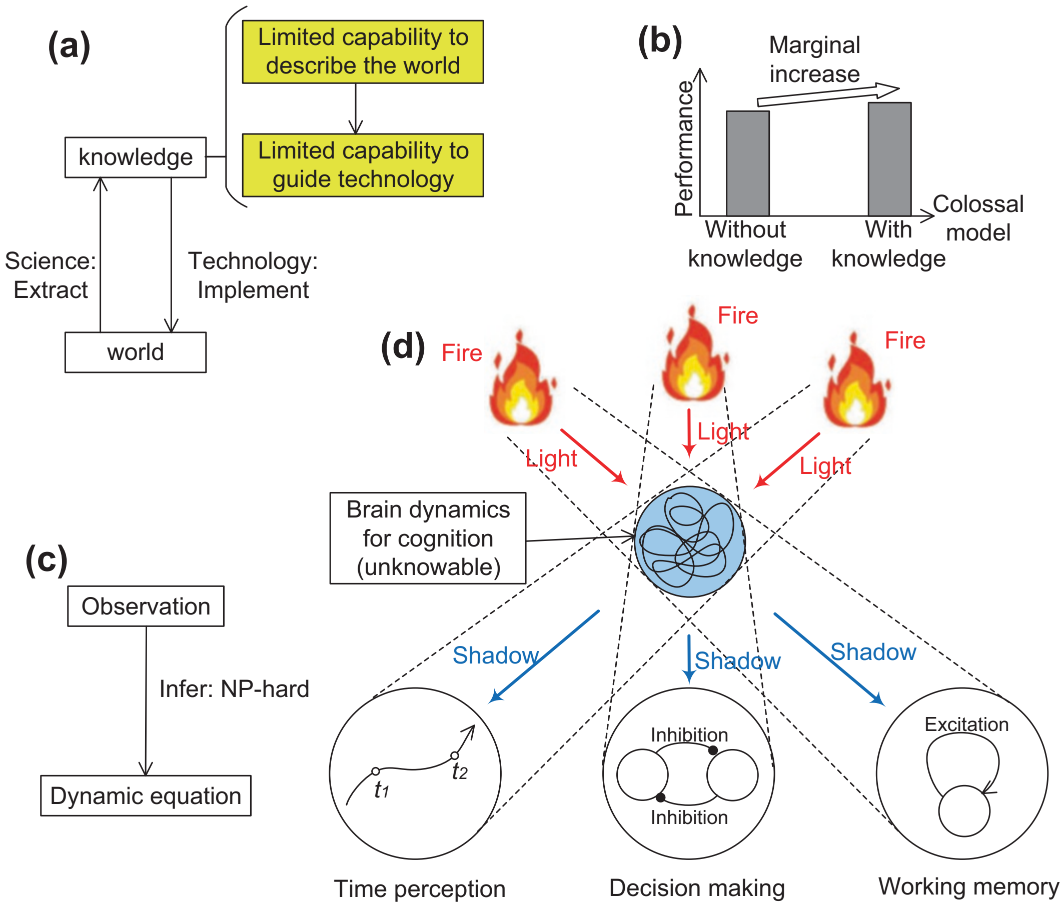

Fig. 6 The unknowability of the world. (**a**) Science extracts knowledge from the world, and technology implements knowledge to change the world. However, if knowledge has limited capability to describe the world, it will also have limited capability to guide the creation of technology (yellow boxes). (**b**) Adding knowledge to colossal models only marginally improves their performance. (**c**) Inferring the dynamic equation from experimental observation is an NP-hard problem. (**d**) A Platonic understanding of cognition. The dynamics of the brain in various tasks (such as the stereotypical trajectory for time perception (Bi & Zhou, 2020a; Karmarkar & Buonomano, 2007), competing dynamics by mutual inhibition for decision-making (Wong & Wang, 2006), and self-excitation for working memory (Lim & Goldman, 2013)) are various shadows of an unknowable object (the brain dynamics for cognition) under fires at different positions. This Platonic viewpoint implies that we will still not understand cognition after studying the dynamics in various tasks, unlike the atomic viewpoint (Fig. 1a), which believes that we will understand cognition after studying each element of cognition

ceived things conform to our spatial and temporal forms of intuition (Immanuel Kant, https://en.wikipedia.org/wiki/Immanuel_Kant). In 1963, Frederic Fitch proposed a logic paradox that asserts that if all truths were knowable, it would follow that all truths are already known (Fitch's paradox of knowability, https://en.wikipedia.org/wiki/Fitch%27s_paradox_of_knowability). Therefore, if we acknowledge that not all truths are already known, we have to acknowledge that not all truths are knowable. Fitch's paradox sets up a fundamental limitation on the capability of experiments: There exists truth that cannot be known using experiments, no matter how advanced the techniques we use.

A recent study in the field of physics provides further evidence for the notion of unknowability. The study demonstrates that identifying the underlying dynamical equation, or physical reality, from any amount of experimental observations is provably NP-hard for both classical and quantum mechanical systems (Cubitt et al., 2012a, b) (Fig. 6c). In simpler terms, if NP/ = P, which is the prevailing belief among computer scientists, identifying the underlying dynamical equation will require an exponentially long amount of time relative to the dimension of the system. Therefore, the dynamical equation would effectively be unknowable if the system has a large dimension.

There have been extensive studies on NP problems using models, such as spin glass, derived from statistical physics (Mézard & Montanari, 2009), which provide insight into the nature of the computational difficulty in solving these problems. The main conclusion is that the computational difficulty is closely related to the (some kind of) correlation between degrees of freedom in the system. To understand this concept, consider a system described by a state vector $x = (x_1, x_2,\ldots, x_n)$. If the different x_is ($i = 1, 2,\ldots, n$) do not interact with each other, we can find the optimal state x_{opt} of the system with respect to a problem by sequentially optimizing each x_i, respectively. However, if the different x_is strongly interact with each other, we may have to adjust a large number of x_is simultaneously during the optimization process, making it more challenging to find x_{opt}.

Basic cognitive studies (Fig. 1) aim to understand the dynamics of the brain underlying cognition by observing the brain's activities when the brain is performing simple tasks. Therefore, basic cognitive studies address the same type of NP-hard problem studied in (Cubitt et al., 2012a, b) that infers dynamics from observation. We have mentioned in the last paragraph that the difficulty of this problem lies in the correlation between different degrees of freedom. Therefore, the atomic philosophy (Fig. 1a), which aims to understand cognition by studying each individual cognitive element (such as perception, memory, and decision-making), is actually unsuitable for guiding cognition research. This is because the coordination between different cognitive elements is vital for performing real-life tasks, so it is important to consider the whole task simultaneously. We have mentioned a good example before (Fig. 4c): The translation of brain activity to language for speech prosthesis achieves better performance when training the neural network to translate one sentence at a time instead of one individual word at a time (Makin et al., 2020; Cogan, 2020; Moses et al., 2019).

Unfortunately, atomism is just the very philosophy that guides basic cognitive studies (including basic timing studies), which is possibly the reason for the difficulty we encounter in understanding cognition. Despite decades of research, we still do not have a complete understanding of how the brain processes time. Results from basic timing studies can sometimes contradict each other (Fig. 5b) and cannot provide guidance for the design of technology. The study of the hippocampus is another example of this issue. While it has been found that the hippocampus transfers memory into the cortex (Goto et al., 2021) and performs inferential reasoning (Barron et al., 2020), the hippocampus encodes place (Sosa & Giocomo, 2021), head directions (Sosa & Giocomo, 2021), time (Eichenbaum, 2014), visual and auditory stimuli (Goto et al., 2021; Turk-Browne, 2019), and abstract knowledge (Nieh et al., 2021), we still do not have a clear understanding of its functional role. In other words, we cannot predict the hippocampus' functional role in a new experimental condition. What is the mechanism of the brain to process time? What is the functional role of the hippocampus? Perhaps, they are essentially unknowable.

How can we make sense of the kaleidoscopic observations in timing and hippocampal studies? In his famous allegory of the cave, Plato likens our understanding of the world to the shadows on the wall of a cave, cast by objects in front of a fire (Allegory of the cave, https://en.wikipedia.org/wiki/Allegory_of_the_cave). Inspired by this allegory, I think the best way to understand the observations in timing or hippocampal studies is to regard the brain dynamics in different experimental conditions as the shadows cast by an object from fires at different positions (Fig. 6d).

The object represents the reality of the neural mechanism of cognition, which is unknowable, but what we can observe are only the dynamics of the brain when performing a specific task. When the fire is at different positions, the projection on the wall is different, just like the kaleidoscopic dynamics of the brain when performing different tasks. This Platonic viewpoint suggests that we may still not understand cognition after studying the dynamics in various tasks, in contrast to the atomic viewpoint (Fig. 1a), which posits that we will understand cognition after studying each element of cognition. Plato encouraged us to walk out of the cave and know the reality of the world through reason. However, inferring the reality from observation is an NP-hard problem (Cubitt et al., 2012a, b), so the reality may be essentially unknowable.

Summary

The results of basic cognitive studies have limited applicability to the development of brain-related technology due to their lack of generalizability. This lack of generalizability may be attributed to the fundamental unknowability of cognition.

Outlook

What can we learn from the understandings above to guide our future research? I give three suggestions, explained in three subsections below (Fig. 7).

Improving Generalizability

As previously mentioned, generalizability is a central aim of science. We want our results to be valid in broader conditions, not just in the specific experimental conditions we investigated (Fig. 7a). Basic cognitive studies (Fig. 1b–e) are typically performed in laboratory experiments, where the experimental conditions are artificially designed and well-controlled, rather than in real-life settings. As previously noted, the lack of generalizability has long been recognized as a shortcoming of laboratory experiments (Brüggemann & Bizer, 2016; Hulstijn, 1997). Therefore, one possible way to improve the generalizability of our results is to extract the features of brain dynamics when subjects are performing real-life tasks, rather than tasks deliberately designed for experiments. Additionally, to

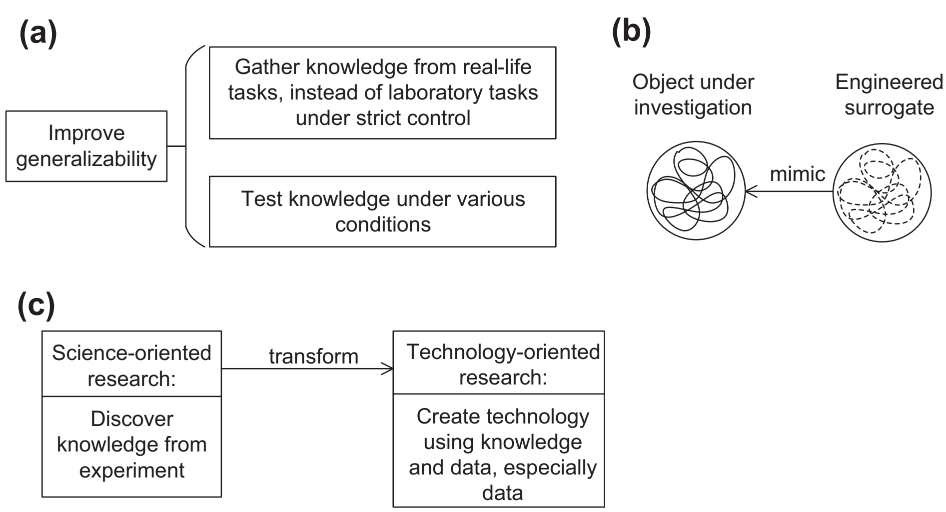

Fig. 7 My suggestions for future research. (**a**) Improve generalizability. (**b**) Engineer an easily accessible surrogate to mimic the object under investigation, so that we can predict the behavior of the object by investigating the surrogate without experimenting on the object. (**c**) Transform our research style from science-oriented, which aims to discover knowledge, to technology-oriented, which aims to create technology

improve the generalizability of a result, it is necessary to verify the result under various conditions.

Making Use of Surrogate

The concept of unknowability suggests that it may be impossible to develop a universal theory that is applicable to every situation. In light of this, how should we proceed with research? One emerging methodology to address the issue of unknowability is the use of surrogates. Instead of attempting to understand the underlying mechanism of a complex system, we create a surrogate of the system that behaves similarly to the original system in situations of interest (Fig. 7b). By observing the behavior of the surrogate, we can predict the behavior of the original system in a new situation. Surrogates are often more accessible than the original system, and this approach can be implemented even when we do not fully understand the mechanism underlying the original system.

Surrogating is the exact idea of neural network models. After observing the input signal I_i ($i = 1, 2, 3, \ldots$), and the output O_i of a system in response to I_i, a deep neural network model can be constructed by training the network to produce the output O_i given the input I_i. The resulting deep network serves as a surrogate for the original system and can predict the output of the system when given a new input signal, I^j, as long as I^j is not significantly different from the set of observed inputs $\{I_i\}_i$. In this approach, the response mechanism of the original system is not explicitly studied; instead, this understanding is encoded in the parameters of the trained neural network. While this knowledge may be difficult to interpret, it can still be used effectively. This type of knowledge is known as "dark knowledge" (Jia, 2019; Hinton et al., 2015), which stands in contrast to "light knowledge" that can be expressed through language and formulas.

Surrogates have been extensively used in brain research. It has been discovered that artificial neural networks, after being trained on a task, exhibit similar dynamics to the brain when performing the same task (Hong et al., 2016; Goldstein et al., 2022; Mante et al., 2013). Therefore, artificial neural networks can be used as surrogates to study the brain, as has been done in the study of timing tasks (Bi & Zhou, 2020a). Additionally, as previously mentioned, neural network models have been utilized as surrogates for epileptic brains to guide clinical surgery (Cao et al., 2022; Sinha et al., 2017).

Furthermore, the concept of surrogating has been applied in fields beyond brain research. For example, self-organized 3-dimensional tissue cultures derived from stem cells, known as organoids, have been used to model various organs, personalize disease treatment, and develop new drugs (Chiaradia & Lancaster, 2020; Kim et al., 2020). Another example is digital twins, which are virtual models designed to accurately represent physical objects and updated in real time with collected data. Digital twins are used to design, manufacture, monitor, and diagnose large equipment such as bridges, aircraft, and power generators (Liu et al., 2021). These examples demonstrate the use of surrogates to investigate and manipulate an easier-to-understand system in order to study the original system, even though the surrogate may also be too complex to be fully understood (e.g., the "dark knowledge" found in artificial neural networks).

Being Technology-Oriented

Science is the process of exploring new knowledge through observation and experiments. Technology is the process of applying scientific knowledge for various purposes. However, the fundamental unknowability of the world presents a limitation on the capability of science to understand the world and guide technology (Fig. 6a). Therefore, in my opinion, future studies should be technology-oriented (Fig. 7c), which has the following two meanings:

1. Instead of being driven by the interest in how nature works, scientists should perform their research with practical applications in mind. A blueprint or at least a rough sketch of how

their findings could be applied to solve practical problems would be beneficial. Without the guidance of technology, scientific results may not be useful in guiding practical applications, as demonstrated by basic timing studies' limited impact on neuroengineering and brain-inspired AI.

2. Technology tends to be created without the guidance of scientific knowledge. There is a growing trend to create technology through human-guided self-organization rather than through the implementation of knowledge via rational design. This shift away from rational design may be due to the unknowability of the world, which renders knowledge increasingly useless in dealing with complex problems. Self-organization is a process in which collective order arises from local interactions between parts of an initially disordered system (Self-organization, https://en.wikipedia.org/wiki/Self-organization; Spontaneous order, https://en.wikipedia.org/wiki/Spontaneous_order). The training of deep artificial neural networks is a self-organization process under human guidance: We adjust the interactions between artificial neurons by adjusting the synaptic weights, rather than directly designing the activity of each neuron, but the collective dynamics of the neural network when performing tasks emerge from these interactions. A good example of the paradigm shift of technology from rational design to human-guided self-organization is natural language processing. The traditional method of translating one language to another was to recognize the grammatical structure of an input sentence and then translate the sentence based on this structure using human-designed rules (Cambria & White, 2014). Today, however, the language translation is based on end-to-end training of neural networks, with the grammatical structure and translation rules automatically and implicitly emerging during training (Goldberg, 2017). As previously mentioned, this automatic and implicit feature extraction by neural networks has also been used to recognize pathological biomarkers in closed-loop deep brain stimulation (Scangos et al., 2021a, b) and translate brain activity into natural language (Makin et al., 2020; Cogan, 2020) (Fig. 4).

How can we guide the self-organization of a complex system to create technology? The current dominating methodology, deep learning, involves adjusting the synaptic weights of a deep network by gradient-based algorithms while fixing the network architecture at the form preassigned by humans (Goodfellow et al., 2016). However, evolutionary algorithms have the potential advantage of allowing for the adjustment of both synaptic weights and network architecture, without requiring human design input (Stanley et al., 2019). In a neural network created by evolutionary algorithms, everything emerges from self-organization, minimizing the interference of human rational design, whose capability is limited due to the unknowability of the world, potentially leading to superior technology (Stanley & Lehman, 2015). Furthermore, human-guided evolution is not only an algorithm that runs on computers but also a practice in laboratories. We create high-yield plants and animals by selective breeding (Selective breeding, https://en.wikipedia.org/wiki/Selective_breeding), and we also discover drugs and functional proteins by directing the evolution of engineered microbes (Davis et al., 2017; Romero & Arnold, 2009). Human-guided evolution, without the need for rational design, may be the ultimate method to create something to our desired end in this unknowable world.

Conclusion

In this paper, I review the main results of basic timing studies and highlight their limited applicability in solving practical problems in the fields of neuroengineering and brain-inspired AI. Basic timing studies extract knowledge from deliberately designed simple tasks, whereas neuroengineering is mainly driven by clinical data and AI is driven by training colossal models using oceans of data collected from the Internet. The limitation of basic timing studies may be due to the lack of generalizability of their results, which stems from the fundamental unknowability of the world, including cognition. The reason for this limitation is also true for, more generally, basic cognitive studies. As a result, I question and criticize the usefulness and prospect of the research protocol of basic cognitive studies (Fig. 1), which involves recording brain activity when the subject is performing deliberately designed experiments to understand the neural mechanism of cognition. I then suggest three ways to guide future research: improving the generalizability of results, considering using surrogates to overcome the unknowability, and performing technology-oriented studies.

The neuroscience problem identified in this paper is part of a larger trend in biology where mass-scale technology, such as multi-omic databases and supercomputing power, is increasingly being used to solve practical problems with AI (Subramanian et al., 2020). The knowledge necessary for AI to solve these problems is not implemented by humans through rational design, but instead emerges self-organizedly during the problem-solving process in a hidden manner. This knowledge is encoded in the AI system, such as in synaptic weights, but is unknowable by humans. We can imagine that in the far future, when AI becomes far more powerful than human intelligence, we may feel hard to understand the logic behind AI's problem-solving even if AI tries to explain it to us. Therefore, the use of hidden knowledge, something we can use but not understand, should be a gradually dominating paradigm in scientific and technological research.

Acknowledgments Z. B. was supported by the National Natural Science Foundation of China (32000694).

References

Allegory of the cave. https://en.wikipedia.org/wiki/Allegory_of_the_cave

Allman, M. J., Teki, S., Griffiths, T. D., & Meck, W. H. (2014). Properties of the internal clock: First- and second-order principles of subjective time. *Annual Review of Psychology, 65*, 743–771.

Angrick, M., et al. (2019). Speech synthesis from ECoG using densely connected 3D convolutional neural networks. *Journal of Neural Engineering, 16*, 036019.

Annervaz, K. M., Chowdhury, S. B. R., & Dukkipati, A. (2018). Learning beyond datasets: Knowledge graph augmented neural networks for natural language processing. *arXiv*, 1802.05930.

Baker, A. (2016). Simplicity. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Stanford University.

Baldassi, C., Braunstein, A., Brunel, N., & Zecchina, R. (2007). Efficient supervised learning in networks with binary synapses. *Proceedings of the National Academy of Sciences of the United States of America, 104*, 11079–11084.

Baldassi, C., Ingrosso, A., Lucibello, C., Saglietti, L., & Zecchina, R. (2015). Subdominant dense clusters allow for simple learning and high computational performance in neural networks with discrete synapses. *Physical Review Letters, 115*, 128101.

Baldwin, J. M. (1893). *Elements of psychology*. Macmillan and Co.

Barron, H. C., et al. (2020). Neuronal computation underlying inferential reasoning in humans and mice. *Cell, 183*, 228–243.

Benabid, A. L., et al. (1994). Acute and long-term effects of subthalamic nucleus stimulation in Parkinson's disease. *Stereotactic and Functional Neurosurgery, 62*, 76–84.

Beniaguev, D., Segev, I., & London, M. (2021). Single cortical neurons as deep artificial neural networks. *Neuron, 109*, 2727–2739.

Bertram, E. H. (2009). Temporal lobe epilepsy: Where do the seizures really begin? *Epilepsy & Behavior, 14*(Suppl 1), 32–37.

Bi, Z., & Zhou, C. (2020a). Understanding the computation of time using neural network models. *Proceedings of the National Academy of Sciences of the United States of America, 117*, 10530–10540.

Bi, Z., & Zhou, C. (2020b). Understanding the computational difficulty of a binary-weight perceptron and the advantage of input sparseness. *Journal of Physics A: Mathematical and Theoretical, 53*, 035002.

Bi, Z., Chen, G., Yang, D., Zhou, Y., & Tian, L. (2022). Evolutionary learning in the brain by heterosynaptic plasticity. *bioRxiv*, 2021.12.14.472260.

Black, J. R. M., & McGranahan, N. (2021). Genetic and non-genetic clonal diversity in cancer evolution. *Nature Reviews. Cancer, 21*, 379–392.

Bommasani, R., et al. (2021). On the opportunities and risks of foundation models. *arXiv*, 2108.07258.

Bouthour, W., et al. (2019). Biomarkers for closed-loop deep brain stimulation in Parkinson disease and beyond. *Nature Reviews. Neurology, 15*, 343–352.

Brown, T. B., et al. (2020). Language models are few-shot learners. *arXiv*, 2005.14165.

Brüggemann, J., & Bizer, K. (2016). Laboratory experiments in innovation research: A methodological overview and a review of the current literature. *Journal of Innovation and Entrepreneurship, 5*, 24.

Bueti, D., & Buonomano, D. V. (2014). Temporal perceptual learning. *Timing and Time Perception, 2*, 261–289.
Buhusi, C. V., & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews. Neuroscience, 6*, 755–765.
Cainelli, E., Mioni, G., Boniver, C., Bisiacchi, P. S., & Vecchi, M. (2019). Time perception in childhood absence epilepsy: Findings from a pilot study. *Epilepsy & Behavior, 99*, 106460.
Cambria, E., & White, B. (2014). Jumping NLP curves: A review of natural language processing research. *IEEE Computational Intelligence Magazine, 9*, 48–57.
Cao, M., et al. (2022). Virtual intracranial EEG signals reconstructed from MEG with potential for epilepsy surgery. *Nature Communications, 13*, 994.
Cheyne, D. O. (2013). MEG studies of sensorimotor rhythms: A review. *Experimental Neurology, 245*, 27–39.
Chiaradia, I., & Lancaster, M. A. (2020). Brain organoids for the study of human neurobiology at the interface of in vitro and in vivo. *Nature Neuroscience, 23*, 1496–1508.
Cichon, J., & Gan, W.-B. (2015). Branch-specific dendritic Ca2+ spikes cause persistent synaptic plasticity. *Nature, 520*, 180–185.
Cogan, G. B. (2020). Translating the brain. *Nature Neuroscience, 23*, 469–472.
Colon-Hernandez, P., Havasi, C., Alonso, J., Huggins, M., & Breazeal, C. (2021). Combining pre-trained language models and structured knowledge. *arXiv*, 2101.12294.
Constantinidis, C., Franowicz, M. N., & Goldman-Rakic, P. S. (2001). Coding specificity in cortical microcircuits: A multiple-electrode analysis of primate prefrontal cortex. *The Journal of Neuroscience, 21*, 3646–3655.
Cook, J. R., et al. (2022). Secondary auditory cortex mediates a sensorimotor mechanism for action timing. *Nature Neuroscience, 25*, 330–344.
Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., & Bengio, Y. (2016). Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or −1. *arXiv*, 1602.02830.
Cubitt, T. S., Eisert, J., & Wolf, M. M. (2012a). Extracting dynamical equations from experimental data is NP hard. *Physical Review Letters, 108*, 120503.
Cubitt, T. S., Eisert, J., & Wolf, M. M. (2012b). The complexity of relating quantum channels to master equations. *Communications in Mathematical Physics, 310*, 383–418.
David Hume. https://en.wikipedia.org/wiki/David_Hume
Davis, A. M., Plowright, A. T., & Valeur, E. (2017). Directing evolution: The next revolution in drug discovery? *Nature Reviews. Drug Discovery, 16*, 681–698.
Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. The MIT Press.
Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv*, 1810.04805.
Eichenbaum, H. (2014). Time cells in the hippocampus: A new dimension for mapping memories. *Nature Reviews. Neuroscience, 15*, 732–744.
Feynman, R. P., Leighton, R. B., & Sands, M. (2011). *The Feynman lectures on physics* (Vol. 1, New Millennium ed.). Basic Books.
Fitch's paradox of knowability. https://en.wikipedia.org/wiki/Fitch%27s_paradox_of_knowability
Frenkel, C. (2021). Sparsity provides a competitive advantage. *Nature Machine Intelligence, 3*, 742–743.
Gámez, J., Mendoza, G., Prado, L., Betancourt, A., & Merchant, H. (2019). The amplitude in periodic neural state trajectories underlies the tempo of rhythmic tapping. *PLoS Biology, 17*, e3000054.
Goldberg, Y. (2017). *Neural network methods in natural language processing*. Morgan & Claypool Publishers.
Goldstein, A., et al. (2022). Shared computational principles for language processing in humans and deep language models. *Nature Neuroscience, 25*, 369–380.
Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. The MIT Press.
Goto, A., et al. (2021). Stepwise synaptic plasticity events drive the early phase of memory consolidation. *Science, 374*, 857–863.
Graupner, M., & Brunel, N. (2010). Mechanisms of induction and maintenance of spike-timing dependent plasticity in biophysical synapse models. *Frontiers in Computational Neuroscience, 4*, 136.
Greyson, B., Fountain, N. B., Derr, L. L., & Broshek, D. K. (2014). Out-of-body experiences associated with seizures. *Frontiers in Human Neuroscience, 8*, 65.
Gu, B.-M., Jurkowski, A. J., Shi, Z., & Meck, W. H. (2016). Bayesian optimization of interval timing and biases in temporal memory as a function of temporal context, feedback, and dopamine levels in young, aged and Parkinson's disease patients. *Timing and Time Perception, 4*, 315–342.
Guo, Q., et al. (2022). A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering, 34*, 3549–3568.
Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron, 95*, 245–258.
Herculano-Houzel, S. (2012). The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proceedings of the National Academy of Sciences of the United States of America, 109*, 10661–10668.
Hetling, J. R. (2008). Comment on 'what is neural engineering?'. *Journal of Neural Engineering, 5*, 360.
Hickok, G. (2012). Computational neuroanatomy of speech production. *Nature Reviews. Neuroscience, 13*, 135–145.
Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv*, 1503.02531.

Hogan, A., et al. (2021). Knowledge graphs. *arXiv*, 2003.02320.
Hong, H., Yamins, D. L. K., Majaj, N. J., & DiCarlo, J. J. (2016). Explicit information for category orthogonal object properties increases along the ventral stream. *Nature Neuroscience, 19*, 613–622.
Hulstijn, J. H. (1997). Second language acquisition research in the laboratory: Possibilities and limitations. *Studies in Second Language Acquisition, 19*, 131–143.
Immanuel Kant. https://en.wikipedia.org/wiki/Immanuel_Kant
Ivry, R. B., & Schlerf, J. E. (2008). Dedicated and intrinsic models of time perception. *Trends in Cognitive Sciences, 12*, 1606–1609.
Janata, P., & Grafton, S. T. (2003). Swinging in the brain: Shared neural substrates for behaviors related to sequencing and music. *Nature Neuroscience, 6*, 682–687.
Ji, D., & Wilson, M. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience, 10*, 100–107.
Jia, W. W. (2019). *Dark knowledge: How machine cognition subverts business and society*. CITIC Press Group.
Jin, D. Z., Fujii, N., & Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proceedings of the National Academy of Sciences of the United States of America, 106*, 19156–19161.
Karmarkar, U. R., & Buonomano, D. V. (2007). Timing in the absence of clocks: Encoding time in neural network states. *Neuron, 53*, 427–438.
Kim, J., Koo, B.-K., & Knoblich, J. A. (2020). Human organoids: Model systems for human biology and medicine. *Nature Reviews. Molecular Cell Biology, 21*, 571–584.
Kirkpatrick, J., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America, 114*, 3521–3526.
Krauss, J. K., et al. (2021). Technology of deep brain stimulation: Current status and future directions. *Nature Reviews. Neurology, 17*, 75–87.
Kuncel, A. M., et al. (2006). Clinical response to varying the stimulus parameters in deep brain stimulation for essential tremor. *Movement Disorders, 21*, 1920–1928.
Lechner, M., et al. (2020). Neural circuit policies enabling auditable autonomy. *Nature Machine Intelligence, 2*, 642–652.
Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews. Neuroscience, 21*, 335–346.
Lim, S., & Goldman, M. S. (2013). Balanced cortical microcircuitry for maintaining information in working memory. *Nature Neuroscience, 16*, 1306–1314.
Limousin, P., et al. (1995). Effect on parkinsonian signs and symptoms of bilateral subthalamic nucleus stimulation. *Lancet, 62*, 91–95.
Liu, M., Fang, S., Dong, H., & Xu, C. (2021). Review of digital twin about concepts, technologies, and industrial applications. *Journal of Manufacturing Systems, 58*, 346–361.
Makin, J. G., Moses, D. A., & Chang, E. F. (2020). Machine translation of cortical activity to text with an encoder- decoder framework. *Nature Neuroscience, 23*, 575–582.
Manning, C. D., Clark, K., Hewitt, J., Khandelwal, U., & Levy, O. (2020). Emergent linguistic structure in artificial neural networks trained by self-supervision. *Proceedings of the National Academy of Sciences of the United States of America, 117*, 30046–30054.
Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature, 503*, 78–84.
Matell, M. S., & Kurti, A. N. (2014). Reinforcement probability modulates temporal memory selection and integration processes. *Acta Psychologica, 147*, 80–91.
Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: Coincidence detection of oscillatory processes. *Cognitive Brain Research, 21*, 139–170.
Matell, M. S., & Valle, R. B. D. (2017). Temporal specificity in Pavlovian-to-instrumental transfer. *Learning & Memory, 25*, 8–20.
Merchant, H., Harrington, D. L., & Meck, W. H. (2013). Neural basis of the perception and estimation of time. *Annual Review of Neuroscience, 36*, 313–336.
Mézard, M., & Montanari, A. (2009). *Information, physics, and computation*. Oxford University Press.
Mita, A., Mushiake, H., Shima, K., Matsuzaka, Y., & Tanji, J. (2009). Interval time coding by neurons in the presupplementary and supplementary motor areas. *Nature Neuroscience, 12*, 502–507.
Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature, 518*, 529–533.
Morita, K., et al. (2020). Clonal evolution of acute myeloid leukemia revealed by high-throughput single-cell genomics. *Nature Communications, 11*, 5327.
Moses, D. A., Leonard, M. K., Makin, J. G., & Chang, E. F. (2019). Real-time decoding of question-and-answer speech dialogue using human cortical activity. *Nature Communications, 10*, 3096.
Nieh, E. H., et al. (2021). Geometry of abstract learned knowledge in the hippocampus. *Nature, 595*, 80–84.
Nielsen, M. A., & Chuang, I. L. (2011). *Quantum computation and quantum information*. Cambridge University Press.
O'Connor, D. H., Wittenberg, G. M., & Wang, S. S.-H. (2005). Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proceedings of the National Academy of Sciences of the United States of America, 102*, 9679–9684.
Okun, M. S. (2012). Deep-brain stimulation for Parkinson's disease. *The New England Journal of Medicine, 367*, 1529–1538.
Oswal, A., Brown, P., & Litvak, V. (2013). Synchronized neural oscillations and the pathophysiology of Parkinson's disease. *Current Opinion in Neurology, 26*, 662–670.

Pasley, B. N., et al. (2012). Reconstructing speech from human auditory cortex. *PLoS Biology, 10*, e1001251.
Patel, A. D. (2003). Language, music, syntax and the brain. *Nature Neuroscience, 6*, 674–681.
Paton, J. J., & Buonomano, D. V. (2018). The neural basis of timing: Distributed mechanisms for diverse functions. *Neuron, 98*, 687–705.
Poewe, W., et al. (2017). Parkinson disease. *Nature Reviews. Disease Primers, 3*, 17013.
Poo, M., et al. (2016). China brain project: Basic neuroscience, brain diseases, and brain-inspired computing. *Neuron, 92*, 591–596.
Pool, J. L. (1954). Psychosurgery in older people. *Journal of the American Geriatrics Society, 2*, 456–466.
Popper, K. (1959). *The logic of scientific discovery*. Hutchinson.
Rakitin, B. C., Gibbon, J., Penney, T. B., & Malapani, C. (1998). Scalar expectancy theory and peak-interval timing in humans. *Journal of Experimental Psychology. Animal Behavior Processes, 24*, 15–33.
Ramesh, A., et al. (2021). Zero-shot text-to-image generation. *arXiv*, 2102.12092.
Rives, A., et al. (2021). Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences of the United States of America, 118*, e2016239118.
Rizzone, M., et al. (2001). Deep brain stimulation of the subthalamic nucleus in Parkinson's disease: Effects of variation in stimulation parameters. *Journal of Neurology, Neurosurgery, and Psychiatry, 71*, 215–219.
Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience, 22*, 9475–9489.
Romero, P., & Arnold, F. (2009). Exploring protein fitness landscapes by directed evolution. *Nature Reviews. Molecular Cell Biology, 16*, 866–876.
Salinas, E., & Sejnowski, T. J. (2001). Gain modulation in the central nervous system: Where behavior, neurophysiology, and computation meet. *The Neuroscientist, 7*, 430–440.
Salinas, E., & Thier, P. (2000). Gain modulation: A major computational principle of the central nervous system. *Neuron, 27*, 15–21.
Scangos, K. W., et al. (2021a). Closed-loop neuromodulation in an individual with treatment-resistant depression. *Nature Medicine, 27*, 1696–1700.
Scangos, K. W., Makhoul, G. S., Sugrue, L. P., Chang, E. F., & Krystal, A. D. (2021b). State-dependent responses to intracranial brain stimulation in a patient with depression. *Nature Medicine, 27*, 229–231.
Schwaller, P., et al. (2021). Mapping the space of chemical reactions using attention-based neural networks. *Nature Machine Intelligence, 3*, 144–152.
Selective breeding. https://en.wikipedia.org/wiki/Selective_breeding
Self-organization. https://en.wikipedia.org/wiki/Self-organization
Shah, S. A., Tinkhauser, G., Chen, C. C., Little, S., & Brown, P. (2018). Parkinsonian tremor detection from subthalamic nucleus local field potentials for closed-loop deep brain stimulation. *Conference Proceedings: Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2320–2324.
Shuler, M. G., & Bear, M. F. (2006). Reward timing in the primary visual cortex. *Science, 311*, 1606–1609.
Siegfried, J., & Lippitz, B. (1994). Bilateral chronic electrostimulation of ventroposterolateral pallidum: A new therapeutic approach for alleviating all parkinsonian symptoms. *Neurosurgery, 35*, 1126–1129.
Sinha, N., et al. (2017). Predicting neurosurgical outcomes in focal epilepsy patients using computational modelling. *Brain, 140*, 319–332.
Sober, E., & Knowles, D. (1991). *Let's Razor Ockham's Razor. Royal Institute of Philosophy Supplements*. Cambridge University Press.
Sosa, M., & Giocomo, L. M. (2021). Navigating for reward. *Nature Reviews. Neuroscience, 22*, 472–487.
Spontaneous order. https://en.wikipedia.org/wiki/Spontaneous_order
Stanley, K. O., & Lehman, J. (2015). *Why greatness cannot be planned: The myth of the objective*. Springer.
Stanley, K. O., Clune, J., Lehman, J., & Miikkulainen, R. (2019). Designing neural networks through neuroevolution. *Nature Machine Intelligence, 1*, 24–35.
Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in nlp. *arXiv*, 1906.02243.
Subramanian, I., Verma, S., Kumar, S., Jere, A., & Anamika, K. (2020). Multi-omics data integration, interpretation, and its application. *Bioinformatics and Biology Insights, 14*, 1177932219899051.
Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. The MIT Press.
Swanton, D. N., & Matell, M. S. (2011). Stimulus compounding in interval timing: The modality-duration relationship of the anchor durations results in qualitatively different response patterns to the compound cue. *Journal of Experimental Psychology. Animal Behavior Processes, 37*, 94–107.
Tan, H., et al. (2019). Decoding voluntary movements and postural tremor based on thalamic LFPs as a basis for closed-loop stimulation for essential tremor. *Brain Stimulation, 12*, 858–867.
Teki, S., Grube, M., Kumar, S., & Griffiths, T. D. (2011). Distinct neural substrates of duration based and beat-based auditory timing. *The Journal of Neuroscience, 31*, 3805–3812.
Turk-Browne, N. B. (2019). The hippocampus as a visual area organized by space and time: A spatiotemporal similarity hypothesis. *Vision Research, 165*, 123–130.
Vaswani, A., et al. (2017). Attention is all you need. *arXiv*, 1706.03762.

Wang, J., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nature Neuroscience, 21*, 102–110.

Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience, 26*, 1314–1328.

Wu, X., et al. (2019). Altered intrinsic brain activity associated with outcome in frontal lobe epilepsy. *Scientific Reports, 9*, 8989.

Zeki, M., & Balci, F. (2019). A simplified model of communication between time cells: Accounting for the linearly increasing timing imprecision. *Frontiers in Computational Neuroscience, 12*, 111.

Zeng, G., Chen, Y., Cui, B., & Yu, S. (2019). Continual learning of context-dependent processing in neural networks. *Nature Machine Intelligence, 1*, 369–372.

Zeng, W., et al. (2021). Pangu-α: Large-scale autoregressive pretrained Chinese language models with auto-parallel computation. *arXiv*, 2104.12369.