

Opinion

Perception in real-time: predicting the present, reconstructing the past

Hinze Hogendoorn  1,2,*,@

We feel that we perceive events in the environment as they unfold in real-time. However, this intuitive view of perception is impossible to implement in the nervous system due to biological constraints such as neural transmission delays. I propose a new way of thinking about real-time perception: at any given moment, instead of representing a single timepoint, perceptual mechanisms represent an entire timeline. On this timeline, predictive mechanisms predict ahead to compensate for delays in incoming sensory input, and reconstruction mechanisms retroactively revise perception when those predictions do not come true. This proposal integrates and extends previous work to address a crucial gap in our understanding of a fundamental aspect of our everyday life: the experience of perceiving the present.

Understanding perception in real-time

Perception is among the most fundamental functions of the nervous system. The function of perception is to organise, identify, and interpret sensory information from the sense organs, in order to represent and understand the environment [1]. For humans, like many other animals, this sensory environment can be highly dynamic, changing rapidly over time. Successfully interacting with objects and agents in the environment therefore requires the perceptual system to accurately represent the sensory environment as it unfolds in real-time.

Understanding how real-time perception is implemented in the human nervous system has turned out to be a surprisingly thorny problem. Our introspective experience appears straightforward: at any given instant, our perception of the environment seems to simply mirror the state of the external environment at that instant (Figure 1). This intuition is so pervasive that the majority of experimental research on perception has largely ignored the fact that both the sensory environment, as well as the neural mechanisms that internally represent that environment, are continually evolving. In a typical laboratory experiment, we might present participants with brief, isolated stimuli, and then record neural responses to each stimulus, or wait until the participant can report some aspect of that stimulus. Such experimental designs are effective because even very brief stimuli (as short as a few milliseconds) can evoke cortical responses that last for a second or more [2], and presenting individual stimuli spaced well apart in this way prevents the responses to different stimuli from running together. However, sensory input outside the laboratory is not so neatly spaced apart. Instead, multiple streams of information from different modalities relentlessly pour into the brain to inform it about the continually evolving sensory environment. Studying perception using individual stimuli therefore both highlights and neatly avoids a fundamental gap in our understanding of the mechanisms underlying perception: how do perceptual mechanisms in the brain cope with the continuous stream of sensory input [3] to generate perception in real-time? Here, I propose a new way to think about how neural mechanisms might enable real-time perception.

Highlights

We feel that we perceive our environment in real-time, despite the constraints imposed by neural transmission delays.

Due to these constraints, the intuitive view of perception in real-time is impossible to implement.

I propose a new way of thinking about real-time perception, in which perceptual mechanisms represent a timeline, rather than a single timepoint.

In this proposal, predictive mechanisms predict ahead to compensate for neural delays, and work in tandem with postdictive mechanisms that revise the timeline as additional sensory information becomes available.

Building on recent theoretical, computational, psychophysical, and functional neuroimaging evidence, this conceptualisation of real-time perception for the first time provides an integrated explanation for how we can experience the present.

¹Melbourne School of Psychological Sciences, The University of Melbourne, Parkville, VIC 3010, Australia
²<http://timinglab.org>

*Correspondence:
hhogendoorn@unimelb.edu.au
(H. Hogendoorn).
©Twitter: @hinzehogendoorn

Three challenges to the intuitive view of real-time perception

Neural delays

There are three key areas where the intuitive view that our perception mirrors the outside world at any given instant falls short. The first is that the transmission and processing of information in the nervous system takes time. During this time, events in the environment continue to unfold, such that sensory information becomes outdated while in transit. In the case of visual motion, for example, a moving object continues moving while sensory information about its position flows through the nervous system. These delays are substantial: it takes several dozen milliseconds for information from the retina to reach visual cortex [4,5], and at least ~120 ms before it is possible to use visual information to initiate voluntary actions [6,7]. In ball sports such as tennis, cricket, or baseball, such delays would correspond to mislocalising the ball several meters behind its true position (Figure 1B). Even a relatively slowly moving object, such as a passing cyclist, would be perceived up to half a meter behind its true position. Humans are nevertheless able to play ball sports and navigate through traffic, and laboratory experiments confirm that humans are remarkably accurate at interacting with dynamic environments, achieving approximately zero lag for even fast-moving objects [8]. So how do perceptual mechanisms compensate for their own delays?

Desynchronisation

The second key challenge to understanding real-time perception is that sensory information becomes desynchronised as it is processed. This is because the time required to process sensory input differs across sensory modalities (e.g., audition precedes vision [9]) and even across features within a modality (e.g., within vision, colour precedes motion [10]). At the neural level, there is enormous variability in the response latency of neurons even within the same visual area [11]. As a result, information from a single sensory event becomes available to perception not at any single moment, but over a range of moments. For the same reason, the most recent information available to perception at any given moment will have originated from different moments in the environment for different features, and therefore do not belong together to the same percept (Figure 1C). This causes a temporal binding problem: as the brain processes the continuous stream of desynchronised sensory input, how does it infer what happened when?

Causality

The final challenge to understanding how perception operates in real-time is the existence of a number of robust perceptual phenomena, in which the second of two sequentially presented stimuli affects the perception of the first. This is problematic because it appears to violate the fundamental rule of physics that causes must precede their consequences; an event cannot cause something to retroactively happen 'back in time'. Nevertheless, there are several known examples where this does appear to happen. In vision, backward masking is one such example, where presenting a second stimulus shortly after an earlier stimulus can render the first stimulus invisible. Conversely, cuing attention to the location of a subthreshold stimulus even after its disappearance can cause the stimulus to be perceived [12]. In the Colour Phi phenomenon (Figure 1D) [13], two differently coloured discs are presented in different positions shortly after one another. This sequence is perceived as a single disc moving smoothly from the first to the second position, changing colour midway. The paradox is that until the second disc is detected, perceptual mechanisms do not have access to its new colour or position, and so should not be able to create a percept of motion and changing colour during the gap. Logically, this would only be possible after the presentation of the second disc. Given that we require the brain to obey the physical law of causality, how do real-time perceptual mechanisms appear to allow sensory information to 'go back in time' to affect the perception of earlier events?

Glossary

Duration channels: a symbolic representation of time in which specific neurons respond to stimuli with specific durations.

Physical time: time as measured on an objective time-scale, which can be used to indicate when physical events (including neural processes) occur. This is distinct from represented time.

Postdiction: a perceptual mechanism in which the perception of a given event is modified by sensory information which is presented after that event. Postdiction appears to violate causality, but does not because it affects the perceptual representation of the past, rather than the past itself.

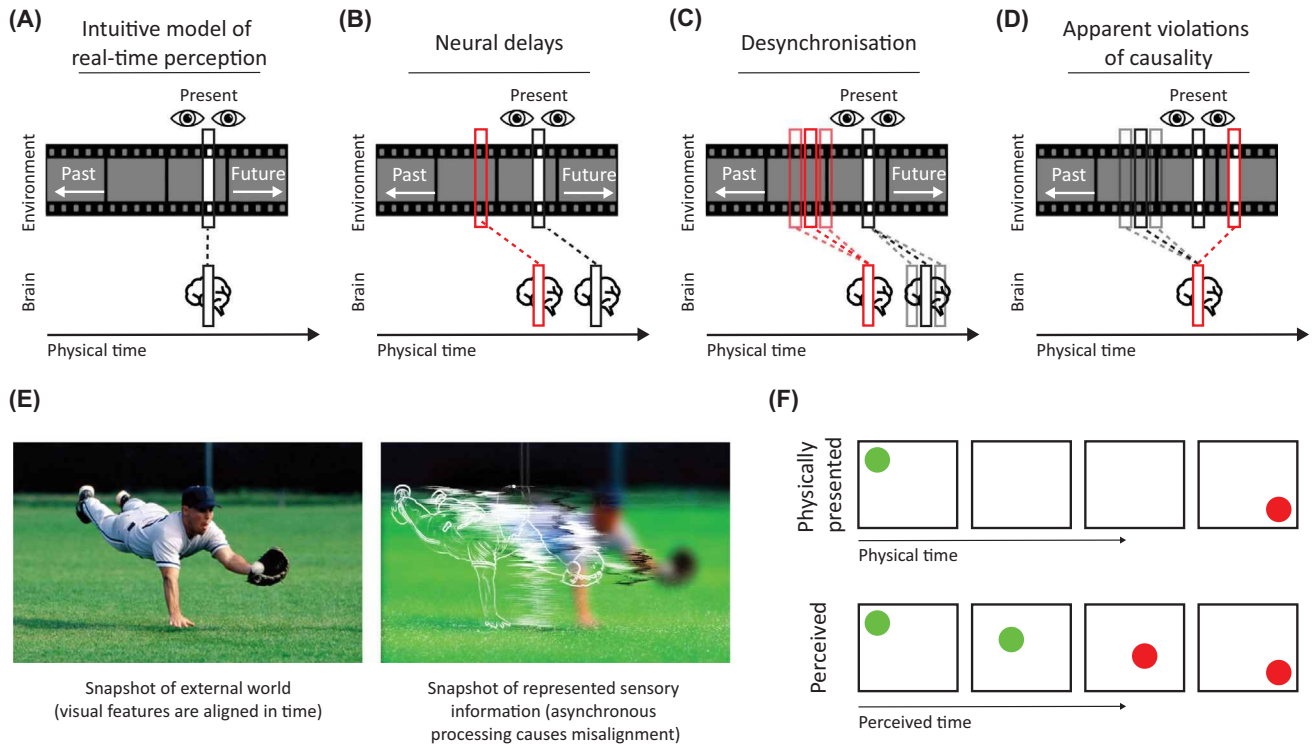
Prediction: a perceptual mechanism by which perceptual information is preactivated even though no sensory information is yet available, using previous sensory information and prior knowledge.

Represented time: time in our subjective experience, corresponding to the time at which we perceive an event to happen, or how long we perceive an event to take. Time in our experience is distinct from the physical timing of the neural mechanisms that create that subjective experience.

Sensory horizon: the point on a perceptual timeline beyond which no sensory information is available, due to neural transmission and processing delays.

Short-term synaptic plasticity (STP): the mechanism by which the relative strength of a neural connection briefly changes as a result of recent activity of that connection.

Symbolic representation: encoding a sensory feature (such as time) as a different dimension of neural activity (for example, as a spatial pattern). The alternative is that the timing of a sensory stimulus is encoded only as the timing of the neural activity signalling that stimulus.



Trends in Cognitive Sciences

Figure 1. The intuitive model of real-time perception, and three challenges that it fails to address. (A) In the intuitive view of perception, we perceive events in the environment as they unfold in the present: the contents of perception are aligned in time with the external world. (B) Due to transmission delays, sensory information about the present moment (black) will only be available in cortex at some point in the future. As a result, perceptual representations that are currently available (red) contain information that is slightly outdated. If this were not somehow compensated, perception would lag behind the present, causing moving objects to be mislocalised behind their true position. (C) Because neural delays are variable across features and modalities, information about the present becomes available over a range of future timepoints (black). As a result, the most recent sensory information available at any given instant (red) does not originate from the same timepoint. (D) Various stimuli can retroactively change the perception of preceding events; a phenomenon called postdiction. In the intuitive model, this means sensory information needs to travel back in time to influence representations that existed earlier, violating the physical law of causality. (E) Desynchronisation of sensory features should cause the perception of different features (e.g., colour, motion, form, etc.) to become misaligned in time, in contrast to our everyday experience. (F) In Colour Phi, the presentation of a red disc after an earlier green disc retroactively creates a percept of motion during the gap, as well as a colour change midway along the trajectory. The intuitive model of real-time perception cannot explain how the disc appears to change colour before the second disc has been presented.

A new conceptualisation of real-time perception

The intuitive model of real-time perception fails to account for important biological constraints and key experimental observations. I propose a new way of thinking about perception in real-time, that addresses the three challenges outlined above.

The key innovation of this proposal is that rather than representing a single timepoint, perceptual mechanisms represent a timeline covering an extent of time. This timeline contains an ongoing best estimate of past, present, and future (Figure 2). It is not a temporal buffer, in which information accumulates before individual timepoints are sequentially perceived; rather, it is an editable sensory narrative, in which sensory interpretations become available to perception as soon as they are detected. Moreover, information on this timeline never becomes definitive, thereby allowing future interpretations to overwrite earlier ones as if those never happened.

This conceptualisation solves a key problem that any account of temporal perception must solve. On the one hand, sensory information rapidly becomes available to guide behaviour (e.g., <150 ms




Figure 2. A new conceptualisation of real-time perception. The key feature of this proposal is that perceptual mechanisms instantaneously represent an entire timeline of events, rather than a single timepoint. Different sensory features (e.g., texture, colour, motion, luminance, etc.) are illustrated here as separate rows on a metaphorical film reel. Sensory input to this timeline is limited by neural delays, imposing a sensory horizon on incoming input that can vary across different features (broken lines). Nevertheless, predictive mechanisms allow sensory information to be represented on this timeline for timepoints beyond this horizon, for which no input is yet available.

[6,7] for saccades). On the other hand, some percepts (such as visual motion) require integration over time, or can be altered by later events (such as Colour Phi). This integration window can extend several hundred ms (e.g., [14,15]). Many other features are perceived only as part of their immediate temporal context (e.g., musical notes in a melody, phonemes in speech, and objects in apparent motion). In an attempt to address this, a recent model of postdictive perception proposed a two-stage model with an unconscious integration window of ~300 ms [15], but this requires that conscious perception lags the external environment by an equivalent duration. Does perception really wait that long before committing to a percept?

In the current proposal, perception does not wait at all. Instead, sensory information becomes available immediately, whilst remaining editable. Although counterintuitive, this provides the best of both worlds: rapid perception is possible when input is available, whilst sensory features that require temporal integration or lead to reinterpretation of earlier input can overwrite the initial percept when they become available. Because there is no separate memory of the initial percept, the timeline always reflects the most coherent and up-to-date interpretation of the occurrence of events.

Representing time

If perception represents a timeline, how does it represent time? In considering how this might be achieved in real-time, it is important to note that the timing of neural activity representing a sensory event is conceptually distinct from the timing of that event, and that both are distinct from the perceived timing of those events. In other words, when an event occurs in the environment, when a sensory representation of that event is activated in the brain, and when the event is perceived to happen are three different questions. Furthermore, whereas the first two can be established on the same objective timeline (**physical time**; see [Glossary](#)), the third can only be answered on a subjective timeline (the person's subjective experience or **represented time**).

It is often implicitly assumed that the perceived timing of events is encoded simply as the timing of the neural activity itself, such that when an event is perceived to happen is determined by when sensory information about that event is processed. However, this encoding strategy of using time to represent time [16–18] cannot explain experimental dissociations of these two timelines [19–21], or how desynchronised sensory information nevertheless leads to an integrated, coherent percept (Figure 1C) [16,22,23].

In the current proposal, the timing of sensory events is instead represented **symbolically**, as the activity of specialised neural systems that encode the relative time of external events just like other features of the event. In this way, when an event is perceived to happen (represented time), is encoded explicitly as a perceptual feature, and is not deduced from (and therefore principally independent of) the time at which sensory information about that event becomes available (physical time). This works in a similar way to a date-stamp on a letter. Due to unknown variations in postage times, the arrival time of a letter is not necessarily an accurate reflection of when it was sent. However, even if two letters are received around the same time, their date-stamps make it possible to ascertain which was sent first. In the same way, even though different sensory features of the same event may be processed at different latencies, and therefore become available to be incorporated into the perceptual timeline at a different physical time, they can still be represented as belonging to the same moment on the internal timeline [24].

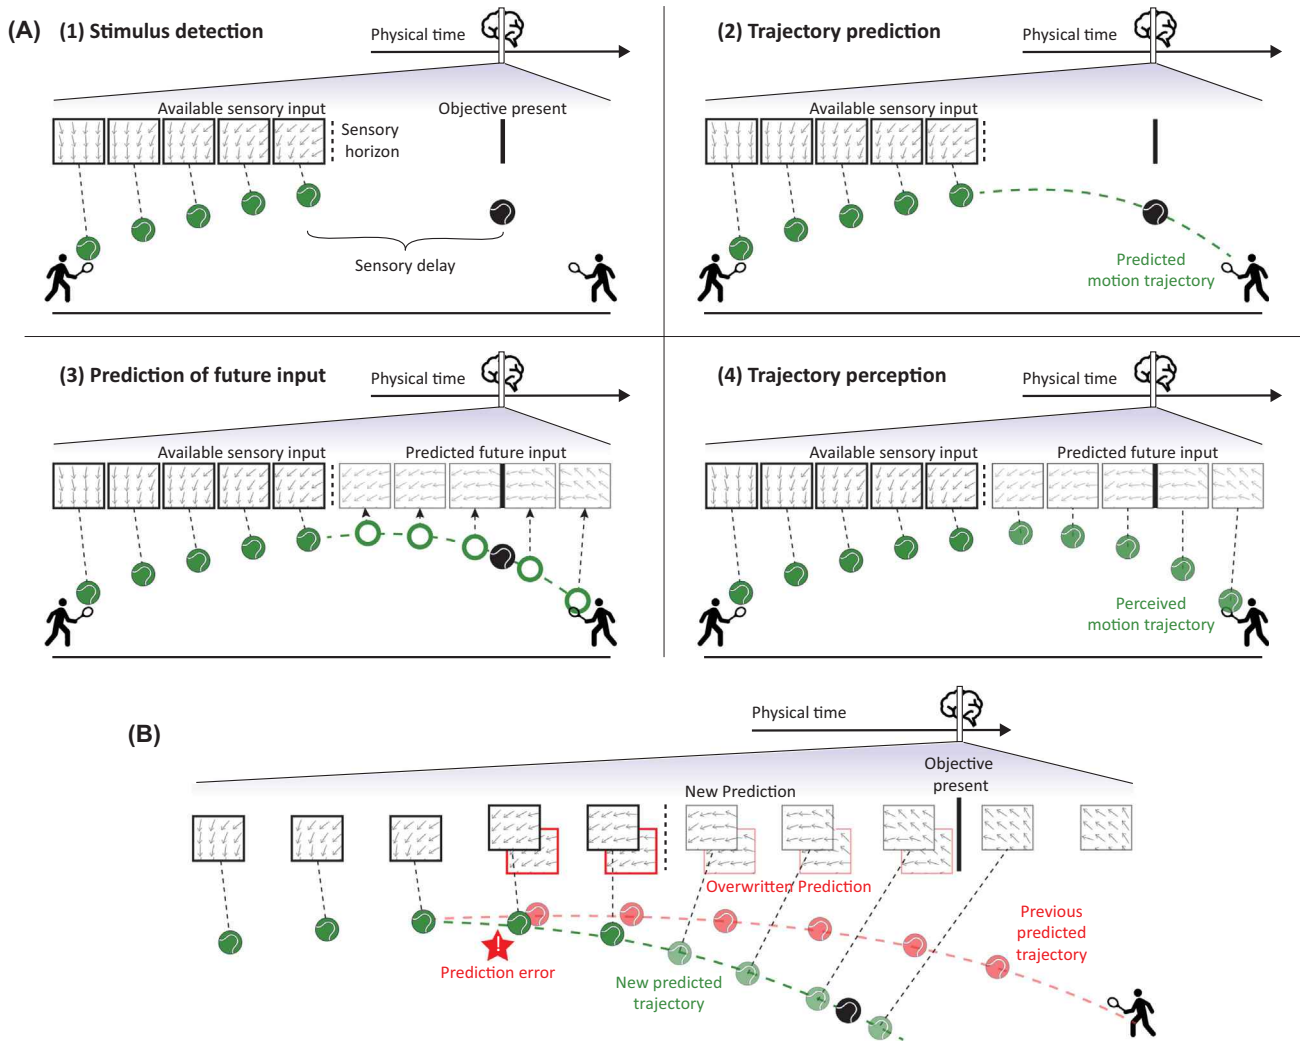
A sensory horizon

Neural delays impose an important constraint on the sensory information that is available to be represented at any given moment. Sensory information about a new event is not available to be incorporated into the timeline until a given amount of time has passed, creating a **sensory horizon**. For this reason, the most recent sensory evidence available to perception always lags slightly behind the present.

I propose that perception compensates for the limitations imposed by the sensory horizon by allowing perceptual representations to be activated by other neural processes, including local computations and feedback from other areas. In particular, **prediction** mechanisms can use available information about past timepoints to estimate the contents of timepoints, for which limited or no sensory information is yet available (Figure 3A and Box 1). In this way, prediction allows a perceptual representation of a particular feature to be formed at a shorter latency than would normally be possible given neural transmission delays.

Furthermore, perceptual representations never become final or definitive. Instead, the timeline is an ongoing best estimate that can change as new evidence becomes available or new interpretations emerge. This means that representations of any timepoint (future, past, or present) can always be updated or revised. Where prediction mechanisms might activate the representation of a future event, **postdiction** processes [14,25] can reconstruct the perceptual past, for example to correct predictions that did not eventuate (Figure 3B and Box 2).

A final implication of the current proposal is that there is no hard natural boundary between perception and memory. Rather, there is a continuum between the two: as perceptual representations become older, they become degraded, compressed, or summarised, gradually becoming experiences of a past event in a way that is typically called episodic memory. This continuum between perception and memory is consistent with previous discussions of consciousness more broadly [26] and postdiction specifically [14], where retroactive revisions of past events are known to take place on timescales ranging from subsecond [14,27] to months or years [28].



Trends In Cognitive Sciences

Figure 3. Complementary roles for prediction and postdiction in updating the perceptual timeline. (A) Due to neural delays, (1) available sensory input about the position of a moving object lags behind the object's physical position (black), but can be used to extrapolate the object's expected trajectory. (2) This allows future sensory input to be predicted (green open circles), such that perceptual representations of timepoints beyond the sensory horizon can be activated. (3) In this way, perception represents a timeline (4), encompassing the ball's entire trajectory, (i.e., past, present, and future). (B) When predictions do not match incoming sensory input, for example because backspin on a tennis ball causes it to deviate from its anticipated trajectory, prediction error results. In this situation, postdiction mechanisms overwrite representations of past or previously predicted events (red), and new predictions are formed (green). The trajectory that was initially perceived is overwritten, and only the new trajectory is perceived and remembered. Importantly, these postdiction mechanisms satisfy causality because they do not affect past representations, but rather overwrite current representations of past events.

Relationship to previous work

Concurrent representation of multiple timepoints

This proposal both integrates and extends previous discussions of perception over time. Firstly, the idea that the nervous system might represent an entire timeline at any given instant builds on a previous proposal that the brain might construct time through trajectory rather than state estimation [17]. If each of the many ways in which the sensory environment can vary constitutes a dimension in a hugely multidimensional space, then the state of the environment at any given instant can be represented by a single point in that state space. As our sensory environment

Box 1. Prediction compensates for neural delays

In the proposed conceptualisation of real-time perception, although sensory information about the environment is not available for timepoints beyond the sensory horizon, these timepoints can nevertheless be activated by prediction. In the case of visual motion, for example, it is possible to extrapolate the previous trajectory of a moving object to predict its real-time position in the present (or even the future), despite neural delays [45,47,88]. Naturally, prediction is only effective at compensating delays when sequences are predictable. If sequences are unpredictable or predictions are violated, then lag is inevitable.

Figure 1A shows an example in which an observer views a bar that rotates smoothly over time. The bar unexpectedly reverses direction between timepoints t_0 and t_1 . Rows in Figure 1B indicate the contents of the perceptual timeline at the four (physical) timepoints t_0 , t_1 , t_2 , and t_3 . Broken squares indicate timepoints for which no sensory input is available at that moment. Asterisks mark represented timepoints that correspond to physical time (i.e., the present).

At time t_0 , predictive mechanisms allow the position of the moving bar in the present to be predicted despite neural delays. However, when the stimulus reverses direction, predictions at time t_1 initially overshoot because the reversal is not yet detected. When incoming sensory information at time t_2 contradicts the predicted representation, the old prediction is discarded (red broken lines), and lag results because no new prediction is yet available. Finally, by time t_3 additional sensory input allows a new prediction to be formed, and the representation ‘catches up’ again.




Figure 1. Prediction and prediction error.

evolves, subsequent moments can be represented as additional points in that same state space. The evolution of our sensory environment over time can then be described by connecting these points into a trajectory: a curve in the system’s state space. I propose that the brain continuously represents and updates that curve, such that the entire perceptual timeline is represented at any given instant.

At the neural level, the implementation of such a process requires that sensory information from different timepoints can be simultaneously represented in a way that still allows them to be

separated. How can the perceptual system represent sensory information from multiple timepoints without having that information blur together? Taking vision as an example modality, one solution to this problem lies in the hierarchical nature of the visual system. Because the transmission of information between layers takes time, at any given instant higher layers represent information from older timepoints than lower layers [29]. Importantly, sensory information is also transformed as it flows through the hierarchy, such that the same information is organised differently at each step along the way [30]. Across its different levels, the hierarchy as a whole therefore automatically represents sensory information from multiple distinct timepoints simultaneously (Figure 4A). A series of recent time-resolved electroencephalogram (EEG) experiments showed this very elegantly by demonstrating that when multiple stimuli are presented in quick succession, several stimuli can be separately decoded from the pattern of brain activity at a single instant [3,31,32].

Symbolic representation of time

This proposal requires that perceptual mechanisms represent time symbolically, as a neural 'time-stamp'. There are at least two ways in which time might be encoded for this purpose. One line of evidence indicates that time perception is encoded by dedicated time-sensitive neurons, such as populations of neurons tuned to specific durations (i.e., **duration channels**), just as there are channels tuned to features such as orientation, velocity, and spatial frequency [33–35]. In this interpretation, the timing of a sensory event is a feature much like its other features, that directly encodes the appropriate position of the represented event on the perceptual timeline. In support of this notion, a number of recent studies have revealed patterns of adaptation to durations that closely mirror the effects of adapting to other features [33–36].

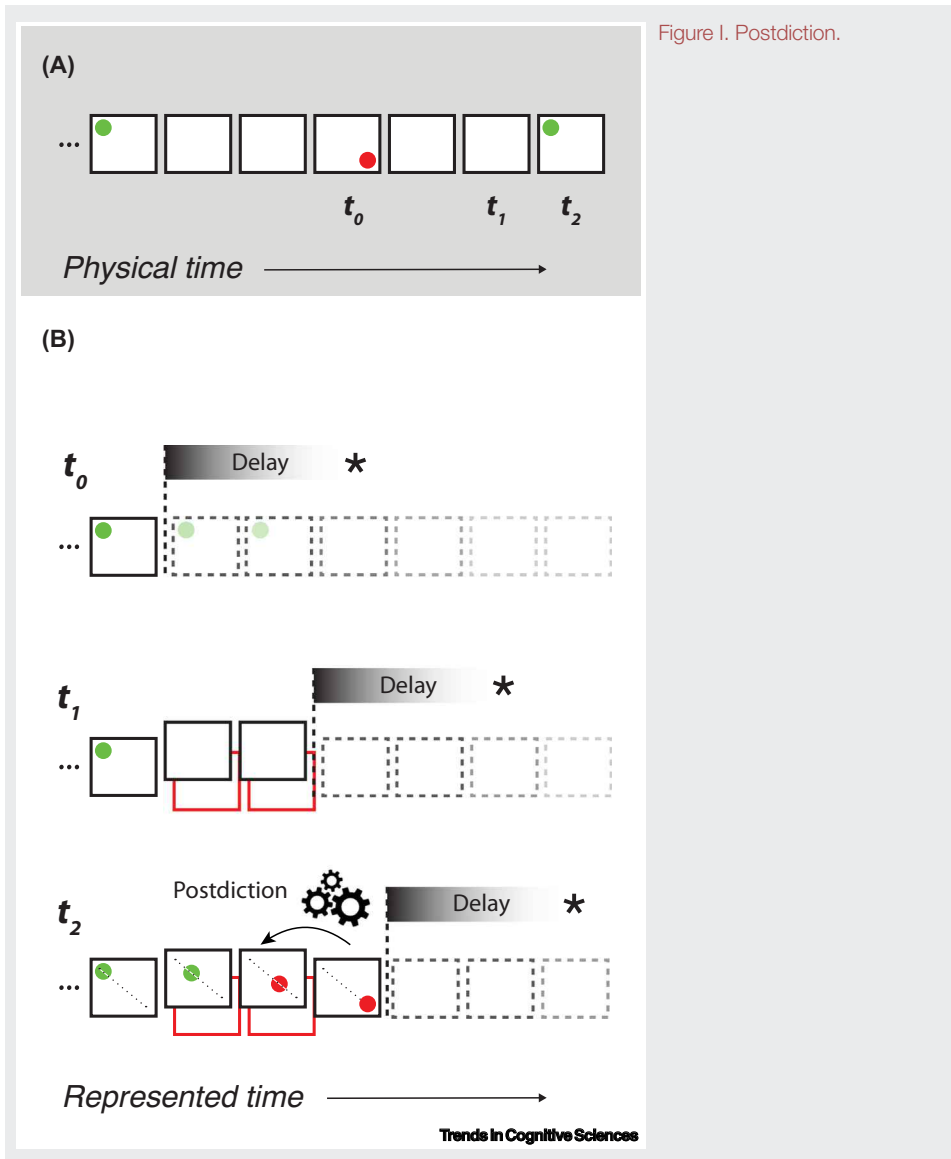
A complementary neural mechanism by which time could be represented is through **short-term synaptic plasticity (STP)**. STP causes recently-active synapses in a neural circuit to subtly change their connectivity. As a result, the exact response of a given neuron to any particular stimulus will depend on its recent history (Figure 4B). At the level of neural circuits, this means that the population response to a stimulus intrinsically encodes the events that preceded the stimulus [37–40]. Simulations in artificial neural circuits show that in this way, the response pattern evoked by a given stimulus in even a small neural population can encode both the identity of previous stimuli as well as how long ago those stimuli were presented [41]. In this way, time (and recent

Box 2. Postdiction reconstructs the perceptual past

A key feature of this proposal is that the perceptual timeline can be updated, revised, reinterpreted, and overwritten as new information (sensory or otherwise) becomes available. This means that the subjective experience of past events can be affected by later events. Importantly, in this account, these postdictive mechanisms do not violate the law of causality because it is the represented past, not the physical past, that is revised.

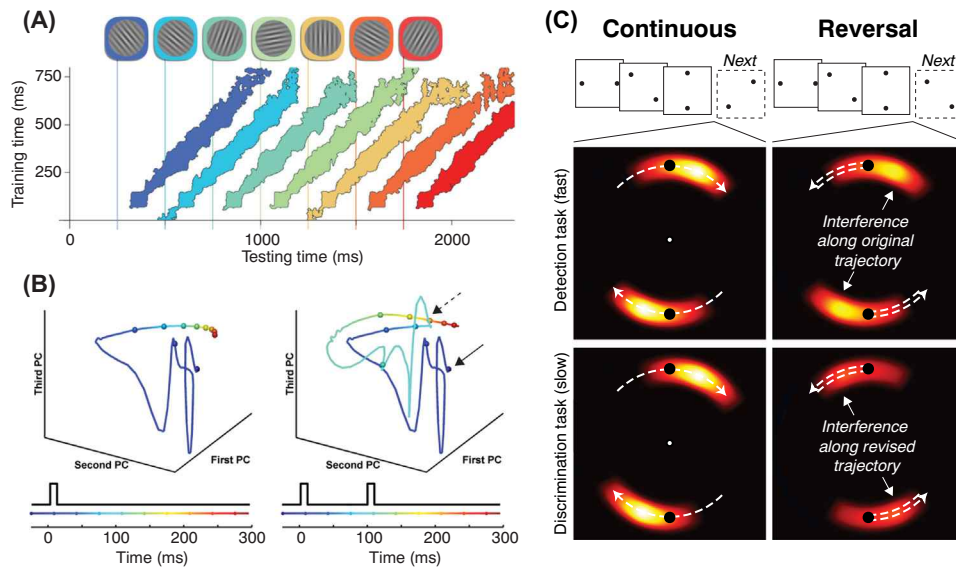
Figure I illustrates how this allows the presentation of a second disc to affect the perception of events leading up that event in the Colour Phi effect [13]. In this phenomenon, observers view two differently coloured discs presented in different positions in quick succession (Figure IA). This creates the percept of a single disc jumping from one position to the other, changing colour midway. As in Box 1, rows in Figure IB indicate the contents of the perceptual timeline at the three (physical) timepoints t_0 , t_1 , and t_2 . Broken squares indicate timepoints for which sensory input is not yet available, and asterisks mark the represented present.

At t_0 , the first available sensory evidence indicates that a disc has been detected. This is represented at the appropriate moment. Future representations may also be activated, depending on prior expectations of the disc's duration. At t_1 , subsequent sensory evidence suggests the disc was an isolated flash. Any earlier prediction is discarded and empty space is represented for the moments following the flash. When the second disc is detected at t_2 , the timeline as a whole is postdictively reinterpreted as a moving disc. The timeline is revised, such that the disc is represented in intervening locations at intermediate moments.



sensory history) are inherently encoded in subtle variations in the spatial pattern of neural responses to the most recent stimulus, without dedicated timing neurons.

Either of these mechanisms would allow perception to encode time-stamps alongside or as part of incoming sensory information. How then might these time-stamps be used to realign sensory input in (represented) time? For one, perception is likely to detect and adapt to temporal offsets in correlated inputs across features, as the sensorimotor system does for voluntary action [42]. Feedback and recurrent connections across and within hierarchical layers allow information from different timepoints to nevertheless be represented concurrently. Indeed, an intriguing EEG study showed that visual luminance sequences continue to replay, or 'echo', in occipital cortex for up to ~1 s [43], providing ample opportunity for coincidence detection and alignment



Trends in Cognitive Sciences

Figure 4. Convergent evidence for the current proposal. (A) A temporal generalisation matrix of EEG data [30] reveals neural representations activated when observers viewed a rapid sequence of oriented visual images. Because individual images evoke waves of neural activity that flow through the visual hierarchy for up to 750 ms, multiple stimuli can be decoded at any given testing time. Reproduced with permission from [3]. (B) Short-term plasticity (STP) causes subtle changes in neural connectivity, such that the response of a neural population to a given stimulus depends on its recent history. Left and right panels show the multidimensional network response to one or two identical brief input pulses, respectively (reduced to three principal component dimensions for illustration). In the right panel, when two identical input pulses are presented, the second pulse (broken arrow) evokes a subtly different neural response to the first pulse (solid arrow), because the network has not yet returned to baseline. In this way, STP causes neural responses to implicitly encode recent sensory history. Reproduced with permission from [40]. (C) Apparent motion reveals both prediction and postdiction. The represented position of an object in apparent motion during the gap between presentations can be inferred from interference on perceptual tasks in intervening positions. When two black discs continue on their clockwise trajectory (left panels), maximal interference is observed ahead of the most recent position, irrespective whether interference is probed with a fast detection task (top) or a slow discrimination task (bottom). When the apparent motion sequence reverses (right panels), at short latencies the object is still represented along its original trajectory (top), but at longer latencies, its represented position is revised such that it now causes interference along the new trajectory (bottom). Reproduced with permission from [77]. Abbreviations: EEG, electroencephalogram; PC, principal component.

with other features. A more recent high-speed fMRI study further revealed that prediction alone is sufficient to preplay sensory sequences in primary visual cortex [44]. Together with mechanisms such as STP, this allows information from different timepoints to be represented concurrently without losing track of where it belongs on the perceptual timeline.

Predicting the present

The proposition that the perceptual system can mitigate the consequences of neural transmission delays through predictive mechanisms has a long history [45–47]. Prediction is essential to accurately interact with a dynamic environment despite these delays, but in principle those delays could be compensated in the motor system [46,48–50], rather than in perception. Nevertheless, for the case of visual motion, convergent evidence from visual illusions (reviewed in [51]), computational modelling [29,52–55], and animal neurophysiology [56,57] suggests that the perceptual system does use prediction to approximate the present [58].

Importantly, recent human neuroimaging work using different sensory and imaging modalities directly supports the notion that predictive mechanisms preactivate neural representations of expected stimuli. For example, the orientation of anticipated gratings can be decoded from the

activation of early visual cortex using both fMRI [59] and magnetoencephalography (MEG) [60]. Likewise, the future position of an apparent motion stimulus can be decoded using both fMRI [44] and EEG [32]. Similar effects have been observed in the auditory system, with MEG revealing preactivation of tonotopic neural representations corresponding to expected but absent tones [61]. Ultra-high field laminar fMRI indicates that such predictions may be driven by feedback to deep cortical layers [62,63]. Importantly, when combined with subsequent afferent sensory information, predictive preactivation both accelerates [64,65] and facilitates [66,67] perception, consistent with theoretical accounts of how predictions [55,68] and memory more generally [69], might integrate with sensory input.

Reconstructing the past

The proposal that the perceptual representation of an event can be modified by the subsequent occurrence of a later event also builds on previous theoretical and experimental work. There are numerous experimental examples of this postdiction phenomenon, including the Flash-Lag [25,70] and Flash-Grab effects [71,72] as well as backward masking and apparent motion [14]. Retro-perception [12] is a further example of later events changing the conscious experience of earlier events [73,74]. Although at first sight paradoxical, postdiction does not entail information going back in time in a way that would violate causality; it merely involves updating a representation of a previous event [14,18,46,75]. Conceptually, prediction and postdiction go hand-in-hand, for the simple reason that predictions do not always eventuate, and failed predictions therefore need to be corrected if the brain is to represent an accurate chronology of events [46].

Importantly, inherent in this reconstruction role of postdiction is that although the original representation is ultimately overwritten, before that happens there is a moment at which the original (predicted) representation is active, however briefly. Although the original representation does not survive to be consciously reported, if it would be possible to probe the perceptual system at precisely the right moment, it should be possible to demonstrate the brief preactivation of that representation. A recent EEG study showed that when an object in apparent motion unexpectedly reverses, this is precisely what is observed [32,76]. When the object changes direction, information about the trajectory change arrives too late to prevent predictive mechanisms from activating perceptual representations along the original trajectory. When sensory information about the new trajectory eventually arrives, the failed prediction is then postdictively overwritten, undoing the preactivation along the original trajectory.

This pattern of neural evidence closely parallels a previous behavioural study [77] that probed the represented position of objects along an apparent motion path using an interference paradigm [78]. Observers carried out a perceptual task on the positions between the possible positions of two discs in apparent motion (Figure 3C). When the discs moved along their original trajectory, maximal interference was observed just ahead of the most recent apparent motion station, as expected. Crucially, when the discs unexpectedly reversed direction, the pattern of interference depended on the perceptual task. With a fast detection task (~ 350 ms), interference was still observed along the original trajectory. However, with a slower discrimination task (~ 620 ms), interference ahead of the last position was lifted, and interference was now observed along the new trajectory (Figure 4C). Together, this pattern of results is consistent with preactivated representations being overwritten by postdictive reconstruction mechanisms when predictions do not eventuate.

Existing models of temporal perception

There is an extensive literature in both philosophy and cognitive neuroscience debating the temporal structure of (conscious) perception [79]. In the first place, there exists a distinction between

three classes of models of the perceived present [80]. In cinematic models, perception takes place at sequential instants, and each percept contains information about a single instant. In retentional models, perception also takes place at sequential instants, but each instant includes information from an extent of time (thereby allowing integration over time). Finally, in extensional models, both perception and its contents are extended in time. My conceptualisation of real-time perception as representing a curve (rather than a point) in a perceptual state space aligns with the class of extensional models [81,82].

An important concept in this context is the specious present: the time period within which percepts are experienced as being 'in the present' [83]. The present proposal extends this concept by challenging the notion that perception is made up of sequential experiences of the present, with previous experiences relegated to memory. Instead, I propose that perception is continuous with memory, such that the (recent) past is, and the predicted future are, experienced alongside the present. This does not cause a 'too-many-percepts' problem [15], because past, present, and future do not overlap as 'now', rather, we perceive a point (or small range, see [Outstanding questions](#)) on this timeline as the present, and also experience the recent past and near future as having just happened and being about to happen, respectively. This also solves the paradox of how perceptual features that are defined over time, such as visual motion or auditory melodies, can be perceived at a single instant [84]. With perception conceived as representing a timeline rather than a single timepoint, it is the object's trajectory in state space that is perceived, rather than a sequence of instantaneous positions [17], thereby solving the paradox.

Another influential debate has centred on the question of whether perception is discrete or continuous [85]. The notion that our apparently continuous stream of perception is made up of a sequence of discrete perceptual frames has been related to neural oscillations [86], but perceptual moments might equally occur 'on demand' [15,87]. Importantly, the current proposal does not require discretization or multistage processing [15,84] to solve the apparent contradiction between the speed of perception and the long integration window of phenomena such as postdiction. This is because under this proposal, postdiction does not impose an upper bound on the latency of conscious perception.

A final question is whether overwritten sensory representations are (briefly) consciously perceived before they are overwritten. When observers are unable to report the original (unrevised) percept in a postdiction paradigm, it is a mistake to conclude that the original input was never perceived. That conclusion assumes that if observers would perceive the original input, they would also have access to a memory trace of that percept by the time they need to report it. However, in most postdiction paradigms, overt responses to a stimulus are executed (at least) several hundreds of milliseconds after the initial event. This means observers' reports of the event are based on a recent representation of that event. This representation is susceptible to postdictive revision [27], and no separate memory trace exists to indicate whether an earlier (now overwritten) representation of that same event was ever perceived. As noted above, the brief existence of the original sensory representation can be demonstrated both behaviourally [77] and neurally [32], but whether it can be said to be briefly consciously perceived (and thus subsequently forgotten) remains an interesting open question (see [Outstanding questions](#)).

Concluding remarks

Despite our everyday introspective experience, the intuitive notion that real-time perception simply reflects the external environment from moment to moment is inconsistent with biological constraints and experimental observations. Here, I address the key challenges caused by neural delays, temporal desynchronisation, and retroactive effects on perception by proposing a new

Outstanding questions

Is there a neural sense in which sensory features processed in distinct brain areas 'all come together', or is conscious perception nothing more than the output of each of the individual areas? Sensory processing is known to be distributed across numerous distinct brain areas, whereas our perceptual experience is generally integrated and coherent. The current proposal presents a mechanism by which different features can be temporally aligned, but how (and if) they combine to produce conscious awareness remains an open question.

What are the precise quantitative computations underlying prediction and postdiction mechanisms, and their integration with sensory input? It seems plausible that some kind of generalised Kalman filter would be involved in generating and updating predictions, as has been successfully modelled for visual motion. However, whether comparable computations operate on other features and modalities remains to be investigated.

How is 'now', the experienced present moment, marked and updated on the perceptual timeline? One point (or small range) is presumably experienced as 'now', but is the position of that point constant relative to the flow of physical time? If the position of this marked present on the timeline is somehow variable, this might lead to distortions of perceived time, such as those that occur around the time of saccades and blinks. It will be interesting to see whether such phenomena can be understood in this framework.

How does the experience of past events change as time passes? How different features of an experience gradually evolve from perceiving now to perceiving the past, to what feels like remembering the past, is an open question. One might wonder if such (presumably gradual) transitions map onto different types of memory, such as iconic, working, and long-term memory, which are also associated with different latencies and capacities.

way to conceptualise real-time perception. The key innovation is that at any given moment, perceptual mechanisms represent an entire timeline, rather than a single timepoint. The proposal is consistent with known properties of the nervous system: neural delays impose a sensory horizon on incoming sensory input to this timeline, but perceptual representations beyond this horizon can nevertheless be preactivated by predictive mechanisms. Perceptual representations of past events can be postdictively updated and revised as new information warrants, without violating principles of causality. Together, predictive and postdictive mechanisms generate a conscious perceptual experience of the present, despite the intrinsic delays in the nervous system.

Acknowledgments

I gratefully acknowledge support from the Australian Research Council (DP180102268 and FT200100246). I am also grateful to Dr Marjolein Kammers and three anonymous reviewers for invaluable comments on earlier versions of this manuscript.

Declaration of interests

No interests are declared.

References

- Schacter, D.L. *et al.*, eds (2010) *Psychology*, 2nd edn Worth Publishers
- Carlson, T. *et al.* (2013) Representational dynamics of object vision: the first 1000 ms. *J. Vis.* 13, 1
- King, J.-R. and Wyart, V. (2021) The human brain encodes a chronicle of visual events at each instant of time thanks to the multiplexing of traveling waves. *J. Neurosci.* 41, 7224–7233
- Maunsell, J.H. and Gibson, J.R. (1992) Visual response latencies in striate cortex of the macaque monkey. *J. Neurophysiol.* 68, 1332–1344
- Maunsell, J.H. *et al.* (1999) Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys. *Vis. Neurosci.* 16, 1–14
- Thorpe, S. *et al.* (1996) Speed of processing in the human visual system. *Nature* 381, 520–522
- Kirchner, H. and Thorpe, S.J. (2006) Ultra-rapid object detection with saccadic eye movements: visual processing speed revisited. *Vis. Res.* 46, 1762–1776
- Smeets, J.B.J. *et al.* (1998) Visuomotor delays when hitting running spiders. In *EWEP 5 - Advances in Perception-Action Coupling* (Bril, B. *et al.*, eds), pp. 36–40, Editions EDK
- Alais, D. and Burr, D. (2003) The “flash-lag” effect occurs in audition and cross-modally. *Curr. Biol.* 13, 59–63
- Arnold, D.H. *et al.* (2001) Asynchronous processing in vision: color leads motion. *Curr. Biol.* 11, 596–600
- Schmolesky, M.T. *et al.* (1998) Signal timing across the macaque visual system. *J. Neurophysiol.* 79, 3272–3278
- Sergent, C. *et al.* (2013) Cueing attention after the stimulus is gone can retrospectively trigger conscious perception. *Curr. Biol.* 23, 150–155
- Kolers, P.A. and von Grünau, M. (1976) Shape and color in apparent motion. *Vis. Res.* 16, 329–335
- Shimojo, S. (2014) Postdiction: its implications on visual awareness, hindsight, and sense of agency. *Front. Psychol.* 5, 196
- Herzog, M.H. *et al.* (2020) All in good time: long-lasting postdictive effects reveal discrete perception. *Trends Cogn. Sci.* 24, 826–837
- Johnston, A. and Nishida, S. (2001) Time perception: brain time or event time? *Curr. Biol.* 11, 427–430
- Grush, R. (2005) Internal models and the construction of time: generalizing from state estimation to trajectory estimation to address temporal features of perception, including temporal illusions. *J. Neural Eng.* 2, S209–S218
- Dennett, D.C. and Kinsbourne, M. (1992) Time and the observer: the where and when of consciousness in the brain. *Behav. Brain Sci.* 15, 183–201
- Johnston, A. (2010) Modulation of time perception by visual adaptation. In *Attention and Time*, pp. 187–200, Oxford University Press
- Johnston, A. *et al.* (2006) Spatially localized distortions of event time. *Curr. Biol.* 16, 472–479
- Hogendoorn, H. *et al.* (2010) Spatially localized time shifts of the perceptual stream. *Front. Psychol.* 1, 181–8
- Moutoussis, K. and Zeki, S. (1997) A direct demonstration of perceptual asynchrony in vision. *Proc. Biol. Sci.* 264, 393–399
- Zeki, S. and Bartels, A. (1999) Toward a theory of visual consciousness. *Conscious. Cogn.* 8, 225–259
- White, P.A. (2020) The perceived present: what is it, and what is it there for? *Psychon. Bull. Rev.* 27, 583–601
- Eagleman, D.M. and Sejnowski, T.J. (2000) Motion integration and postdiction in visual awareness. *Science* 287, 2036–2038
- Edelman, G. (2006) Consciousness: the remembered present. *Ann. N. Y. Acad. Sci.* 929, 111–122
- Intraub, H. and Dickinson, C.A. (2008) False memory 1/20th of a second later. *Psychol. Sci.* 19, 1007–1014
- Loftus, E.F. and Palmer, J.C. (1974) Reconstruction of automobile destruction: an example of the interaction between language and memory. *J. Verbal Learn. Verbal Behav.* 13, 585–589
- Hogendoorn, H. and Burkitt, A.N. (2019) Predictive coding with neural transmission delays: a real-time temporal alignment hypothesis. *eNeuro* 6 ENEURO.0412-0418.2019
- King, J.-R. and Dehaene, S. (2014) Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn. Sci.* 18, 203–210
- Grootswagers, T. *et al.* (2019) The representational dynamics of visual objects in rapid serial visual processing streams. *Neuroimage* 188, 668–679
- Blom, T. *et al.* (2020) Predictions drive neural representations of visual events ahead of incoming sensory information. *Proc. Natl. Acad. Sci. U. S. A.* 117, 7510–7515
- Heron, J. *et al.* (2012) Duration channels mediate human time perception. *Proc. R. Soc. B Biol. Sci.* 279, 690–698
- Maarseveen, J. *et al.* (2019) The duration aftereffect does not reflect adaptation to perceived duration. *PLoS One* 14, e0213163
- Maarseveen, J. *et al.* (2017) An investigation of the spatial selectivity of the duration after-effect. *Vis. Res.* 130, 67–75
- Maarseveen, J. *et al.* (2018) Attention gates the selective encoding of duration. *Sci. Rep.* 8, 2522
- Buonomano, D. and Merzenich, M. (1995) Temporal information transformed into a spatial code by a neural network with realistic properties. *Science* 267, 1028–1030
- Motanis, H. *et al.* (2018) Short-term synaptic plasticity as a mechanism for sensory timing. *Trends Neurosci.* 41, 701–711

39. Buonomano, D.V. and Maass, W. (2009) State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* 10, 113–125
40. Karmarkar, U.R. and Buonomano, D.V. (2007) Timing in the absence of clocks: encoding time in neural network states. *Neuron* 53, 427–438
41. Paton, J.J. and Buonomano, D.V. (2018) The neural basis of timing: distributed mechanisms for diverse functions. *Neuron* 98, 687–705
42. Haggard, P. *et al.* (2002) Voluntary action and conscious awareness. *Nat. Neurosci.* 5, 382–385
43. VanRullen, R. and Macdonald, J.S.P. (2012) Perceptual echoes at 10 Hz in the human brain. *Curr. Biol.* 22, 995–999
44. Ekman, M. *et al.* (2017) Time-compressed preplay of anticipated events in human primary visual cortex. *Nat. Commun.* 8, 15276
45. Nijhawan, R. (2008) Visual prediction: psychophysics and neurophysiology of compensation for time delays. *Behav. Brain Sci.* 31, 179–239
46. Eagleman, D.M. (2008) Prediction and postdiction: two frameworks with the goal of delay compensation. *Behav. Brain Sci.* 31, 205–206
47. Nijhawan, R. (1994) Motion extrapolation in catching. *Nature* 370, 256–257
48. Nijhawan, R. and Wu, S. (2009) Compensating time delays with neural predictions: are predictions sensory or motor? *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* 367, 1063–1078
49. Nijhawan, R. and Kirschfeld, K. (2003) Analogous mechanisms compensate for neural delays in the sensory and the motor pathways. *Curr. Biol.* 13, 749–753
50. Kerzel, D. and Gegenfurtner, K.R. (2003) Neuronal processing delays are compensated in the sensorimotor branch of the visual system. *Curr. Biol.* 13, 1975–1978
51. Hogendoorn, H. (2020) Motion extrapolation in visual processing: lessons from 25 years of flash-lag Debate. *J. Neurosci.* 40, 5698–5705
52. Kwon, O.-S. *et al.* (2015) Unifying account of visual motion and position perception. *Proc. Natl. Acad. Sci. U. S. A.* 112, 8142–8147
53. Burkitt, A.N. and Hogendoorn, H. (2021) Predictive visual motion extrapolation emerges spontaneously and without supervision at each layer of a hierarchical neural network with spike-timing-dependent plasticity. *J. Neurosci.* 41, 4428–4438
54. Khoei, M.A. *et al.* (2017) The flash-lag effect as a motion-based predictive shift. *PLoS Comput. Biol.* 13, e1005068
55. Feuerriegel, D. *et al.* (2021) Predictive activation of sensory representations as a source of evidence in perceptual decision-making. *Cortex* 136, 140–146
56. Schwartz, G. *et al.* (2007) Synchronized firing among retinal ganglion cells signals motion reversal. *Neuron* 55, 958–969
57. Berry, M.J. *et al.* (1999) Anticipation of moving stimuli by the retina. *Nature* 398, 334–338
58. White, P.A. (2018) Is the perceived present a predictive model of the objective present? *Vis. Cogn.* 26, 624–654
59. Kok, P. *et al.* (2014) Prior expectations evoke stimulus templates in the primary visual cortex. *J. Cogn. Neurosci.* 26, 1546–1554
60. Kok, P. *et al.* (2017) Prior expectations induce prestimulus sensory templates. *Proc. Natl. Acad. Sci. U. S. A.* 114, 10473–10478
61. Demarchi, G. *et al.* (2019) Automatic and feature-specific prediction-related neural activity in the human auditory system. *Nat. Commun.* 10, 3440
62. Kok, P. *et al.* (2016) Selective activation of the deep layers of the human primary visual cortex by top-down feedback. *Curr. Biol.* 26, 371–376
63. Aitken, F. *et al.* (2020) Prior expectations evoke stimulus-specific activity in the deep layers of the primary visual cortex. *PLoS Biol.* 18, e3001023
64. Hogendoorn, H. and Burkitt, A.N. (2018) Predictive coding of visual object position ahead of moving objects revealed by time-resolved EEG decoding. *Neuroimage* 171, 55–61
65. Pinto, Y. *et al.* (2015) Expectations accelerate entry of visual stimuli into awareness. *J. Vis.* 15, 13–15
66. Aitken, F. *et al.* (2020) Prior expectations of motion direction modulate early sensory processing. *J. Neurosci.* 40, 6389–6397
67. Roach, N.W. *et al.* (2011) Visual motion induces a forward prediction of spatial pattern. *Curr. Biol.* 21, 740–745
68. de Lange, F.P. *et al.* (2018) How do expectations shape perception? *Trends Cogn. Sci.* 22, 764–779
69. Nobre, A.C. and Stokes, M.G. (2019) Remembering experience: a hierarchy of time-scales for proactive attention. *Neuron* 104, 132–146
70. Rao, R.P.N. *et al.* (2001) Optimal smoothing in visual motion perception. *Neural Comput.* 13, 1243–1253
71. Blom, T. *et al.* (2019) When predictions fail: correction for extrapolation in the flash-grab effect. *J. Vis.* 19, 3–11
72. Cavanagh, P. and Anstis, S.M. (2013) The flash grab effect. *Vis. Res.* 91, 8–20
73. Rinsky-Robert, D. *et al.* (2019) Retrospective auditory cues can improve detection of near-threshold visual targets. *Sci. Rep.* 9 18966–11
74. Thibault, L. *et al.* (2016) Retrospective attention gates discrete conscious access to past sensory stimuli. *PLoS One* 11, e0148504
75. Sergent, C. (2018) The offline stream of conscious representations. *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* 373, 20170349
76. Blom, T. *et al.* (2021) The time-course of prediction formation and revision in human visual motion processing. *Cortex* 138, 191–202
77. Hogendoorn, H. *et al.* (2008) Interpolation and extrapolation on the path of apparent motion. *Vis. Res.* 48, 872–881
78. Yantis, S. and Nakama, T. (1998) Visual interactions in the path of apparent motion. *Nat. Neurosci.* 1, 508–512
79. Kent, L. and Wittmann, M. (2021) Time consciousness: the missing link in theories of consciousness. *Neurosci. Conscious.* 2021, niab011
80. Dorato, M. and Wittmann, M. (2020) The phenomenology and cognitive neuroscience of experienced temporality. *Phenomenol. Cogn. Sci.* 19, 747–771
81. Hoerl, C. (2014) Time and the domain of consciousness. *Ann. N. Y. Acad. Sci.* 1326, 90–96
82. Dainton, B. (2011) Time, passage, and immediate experience. In *Oxford Handbook of Philosophy of Time* (Callendar, C., ed.), pp. 382, Oxford University Press
83. Andersen, Holly K. and Grush, Rick (2009) A brief history of time-consciousness: historical precursors to James and Husserl. *J. Hist. Philos.* 47, 277–307
84. Herzog, M.H. *et al.* (2016) Time slices: what is the duration of a percept? *PLoS Biol.* 14, e1002433
85. VanRullen, R. and Koch, C. (2003) Is perception discrete or continuous? *Trends Cogn. Sci.* 7, 207–213
86. VanRullen, R. (2016) Perceptual cycles. *Trends Cogn. Sci.* 20, 723–735
87. Brenner, E. and Smeets, J.B. (2000) Motion extrapolation is not responsible for the flash-lag effect. *Vis. Res.* 40, 1645–1648
88. Nijhawan, R. (2002) Neural delays, visual motion, and the flash-lag effect. *Trends Cogn. Sci.* 6, 387