Trends in
Cognitive Sciences

CellPress

Review

The causal structure and computational value of narratives

Janice Chen [1,*] and Aaron M. Bornstein [2,3]

Many human behavioral and brain imaging studies have used narratively structured stimuli (e.g., written, audio, or audiovisual stories) to better emulate real-world experience in the laboratory. However, narratives are a special class of real-world experience, largely defined by their causal connections across time. Much contemporary neuroscience research does not consider this key property. We review behavioral and neuroscientific work that speaks to how causal structure shapes comprehension of and memory for narratives. We further draw connections between this work and reinforcement learning, highlighting how narratives help link causes to outcomes in complex environments. By incorporating the plausibility of causal connections between classes of actions and outcomes, reinforcement learning models may become more ecologically valid, while simultaneously elucidating the value of narratives.

Highlights

Naturalistic stimuli used in modern human neuroscience research are often narratives, but there has been limited discussion of a core property of these narratives: causal structure.

Causal structure of narratives has a powerful impact on episodic memory, and emerging evidence suggests that causal variables play a role in the activity of brain systems during narrative encoding and retrieval.

Consideration of how humans identify and use causal information in narratives may inform studies and computational models of the neural mechanisms that enable organisms to link actions and outcomes.

Narrative schema may be especially valuable for improving value-guided learning and decision-making by enabling efficient credit assignment in complex, high-dimensional state spaces like those encountered in natural environments.

Causal structure in narrative naturalistic stimuli

The use of naturalistic stimuli to study human brain function is motivated by the idea that they resemble real-world experience in terms of complexity, content, and statistics. This theoretically enables investigations of neural and mental phenomena that are unavailable when using more controlled and simplified stimuli. One line of questions, mainly within vision science, has focused on how using artificial displays versus natural images or videos affects neuronal response properties in both humans and non-human animals [1–3]. Related lines of research have focused on higher-level cognitive processes in the brain, such as semantic representations [4,5], memory [6–9], attention [10], and affect [11,12]. Such studies require stimuli with cognitive complexity in addition to dynamic visual and auditory features, and thus have tended to draw from entertainment media, such as cinematic movies, television, podcasts, and written fiction. These naturalistic stimuli are typically narratives, whether they are audiovisual, audio-only, or visual-only. A narrative may be defined, broadly, as a series of events [13], and recent cognitive neuroscience studies have been inspired by the classic psychological literature on event models [14,15]. However, while numerous studies have examined neural correlates of event perception and event segmentation, far fewer have addressed a key property of narratives that featured prominently in that same classic literature: narratives are networks of causally connected events, often forming chains of causal links, and these causal connections have a significant impact on comprehension and memory.

Recognizing the critical role of causal structure in narrative stimuli, and of causal reasoning in the processing of such stimuli, can enhance understanding of prior results and guide the design of future experiments. It may also be possible for insights from studies of narratives to motivate questions in research on models of learning. For example, an important problem in reinforcement learning is how to identify the specific prior actions or events that caused the current reward or punishment (**credit assignment**; see Glossary); this problem becomes especially challenging when the causal event occurred many events ago, with counterfactual alternatives along the

[1]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD, USA
[2]Department of Cognitive Sciences, University of California, Irvine, CA, USA
[3]Center for the Neurobiology of Learning and Memory, University of California, Irvine, CA, USA

*Correspondence:
janice@jhu.edu (J. Chen).

https://doi.org/10.1016/j.tics.2024.04.003

way. In narrative comprehension, causal connections between events are assessed by the reader based on their world knowledge of what types of actions may plausibly cause certain outcomes; for example, if a king falls ill, a poisoned meal is a plausible cause while the arrival of a letter is less so. Reinforcement learning models could achieve greater ecological validity by incorporating the plausibility of causal connections between classes of actions and outcomes, as organisms likely do when they make decisions in real-world scenarios. Beyond individual action-outcome pairs, the global causal structure of narratives is often familiar (scripts [16,17]), which provides information about the likelihood of causal connections; for example, if a wife demands divorce when her husband forgets to take out the trash, observers know that this minor incident is not the real cause, but that many prior events probably jointly contributed to her current decision. Models of learning could investigate how transferrable causal structure in these multievent sequences is acquired across exposures where details may differ between instances. In return, a reinforcement learning perspective may be useful in elucidating the evolutionary reasons for certain properties of narratives and why people seem to value them highly. Accurate credit assignment is crucial to efficiently and effectively navigating a complex world, and narratives may constitute compact descriptions of the set of causal links that likely contributed to a given outcome (reward or punishment), extracted via inference and reasoning across past events and compiled into a form which is easy to encode into long-term memory, as well as communicable to others.

In this review, we survey behavioral and neuroscientific studies which bear on the question of how causal structure in narratives shapes memory, highlighting investigations of neural systems which support episodic memory and mental simulation. In parallel, we draw links to work on causal attribution in studies of reinforcement learning and suggest ways in which the two literatures may inform each other, with the goal of achieving a deeper understanding of the value of narratives.

Behavioral investigations of how causal structure in narratives impacts memory

Narrative comprehension was a highly active area of research beginning in the 1970s. These early studies used written narratives to investigate how readers build mental representations of events and what factors determine later memory. A key idea in this literature was that readers must go beyond the literal or surface features of text to generate meaningful representations of the events described in the story, and that doing so depends upon drawing on world knowledge; both surface form and meaning can be remembered, but they are distinct [18–21]. Researchers also examined how representation formation during narrative comprehension is guided by canonical schemas and meaningful relations among text elements (e.g., clauses or events), with more coherent representations yielding better memory performance [22,23].

Several studies used networks to model mental representations of narratives. These studies found that causal relations between text elements or events (the nodes of the networks) are of critical importance in determining comprehension and memory. Narratives were typically divided into events, for example, sentences or short passages (Figure 1A), and causal links between events were classified by experts (e.g., the researchers) according to a variety of logical criteria, such as temporal priority, operativity, necessity, and sufficiency [23–25]. In recent studies, causal links between events are more often identified by asking naive participants for their layman's opinion of whether a causal relationship is present [26] (Figure 1B). The causal structure of a narrative can then be visualized as an event-by-event matrix, either directed or undirected (Figure 1C), and as a network (Figure 1D). Some key findings from this literature were that creating a narrative from otherwise unrelated items dramatically improves recollection [27], the number of causal connections per event predicts judgments of importance and improved memory [23,24,28,29], and causal chains help organize event memory and determine what will be retained [30,31]. Studies

Glossary

Credit assignment: the problem of selecting which among many preceding actions or states should be assigned value crediting it for being predictive of a given experienced reward.

Curse of dimensionality: broadly, refers to the many different issues that arise when working with large, high-dimensional data that do not appear in lower dimensions. For the purposes of this article, we use this term to refer to the difficulty of credit assignment in large, naturalistic environments with many dynamic features and action sequences.

Discounting: the degree to which future rewards are valued less than present rewards. This factor can be justified by uncertainty about the receipt of future rewards, or by their lesser utility with respect to present goals.

Reward prediction errors (RPEs): upon entry into a state, the difference between the value expected and that observed. Critically, in reinforcement learning models, the value signal is the (discounted) sum of all future rewards that could be expected to follow, until the termination of the task episode.

State space: an abstraction used in reinforcement learning to distinguish configurations or situations possible in a given environment. States can be defined at multiple levels of abstraction. An ice cream store may be described as a single state, or it may be multiple states that vary with what is in stock. The level of abstraction is chosen to match the goals of the learning agent – more detailed state spaces can lead to more nuanced policies, at the cost of greater representational and computational complexity.

Successor representation: a compact representation of the likelihood that trajectories beginning from each possible state in a given environment will eventually occupy each of the other possible states; these can be computed either at the full horizon of the task (e.g., extending to the terminal state), or at multiple intermediate timescales, giving n-step successor representations.

Value of information: the increase in future expected value that results from obtaining a given piece of information. This quantity can result from an increase in the likelihood of a given future reward, or an increase in certainty about the delivery of rewards.

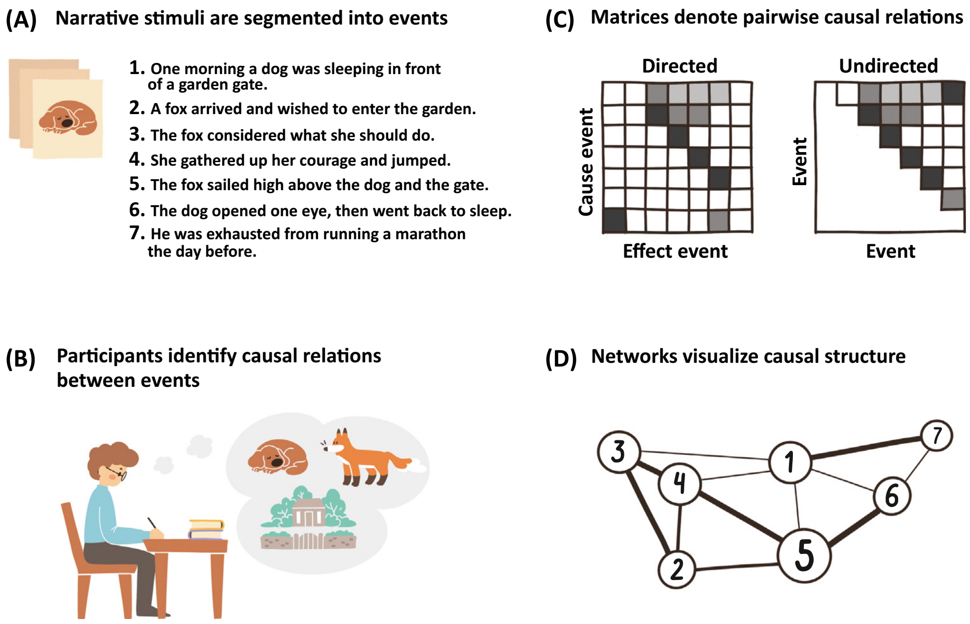

1. One morning a dog was sleeping in front of a garden gate.
2. A fox arrived and wished to enter the garden.
3. The fox considered what she should do.
4. She gathered up her courage and jumped.
5. The fox sailed high above the dog and the gate.
6. The dog opened one eye, then went back to sleep.
7. He was exhausted from running a marathon the day before.

Trends in Cognitive Sciences

Figure 1. Creating and depicting narrative causal networks. (A) A narrative stimulus, such as a written story or a movie, is first divided into events. In a written story, a sentence or a short passage typically corresponds to one event; in a movie, events may correspond to a short scene. In both cases, event boundaries are often defined by collecting judgments of event boundaries from a participant group [80]. (B) Judgments are collected from participants about the causal relationships between pairs of events [26,58]. (C) Pairwise relations between events can be plotted as directed or undirected matrices; directed matrices contain information about the direction of causality (which event caused which other event), while undirected matrices contain information about whether a link exists or not regardless of direction. The value in each cell indicates the strength of the connection; for example, the proportion of the participants who identified that event pair as being causally linked. (D) Narrative causal event matrices can also be visualized as networks (undirected version shown here), wherein each event from the narrative is represented as a node. Node size indicates weighted degree centrality, that is, the number and strength of connections, and edge thickness reflects the connection strength between events [26]. (Image credit: Anna Kuzmina, @kiki_arts).

have also shown that recall is influenced by causal links more than temporal sequence [32], and enhanced for sentences that create causal breaks in a narrative [33].

Do causal connections in narratives imply causal reasoning during narrative comprehension? Theories of causal reasoning often define causal relations between events according to how one event's existence affects the other [26]. For example, 'if A occurs, B always occurs' we would say that 'A causes B'; 'if A occurs, B never occurs' we would say that 'A prevents B'. For simple physical, social, or biological scenarios, such relations can be studied in a highly controlled manner; for example, asking participants to judge whether one ball caused another to miss a target [34], or what factors caused illness in a fictional world [35]. One possibility is that causal reasoning during narrative comprehension draws on the same mechanisms modeled in simpler tasks, combining them as needed. However, the complexity of reasoning needed to comprehend popular culture stories and movies seems to exceed the scope of current models. For example, an event in a narrative commonly has multiple candidate causes (Figure 2), a challenging scenario for formal models of reasoning. Additionally, processing a narrative might not require acts of causal reasoning at every event that is causally linked to a prior one; for example, if in Event 1 a wizard sends a letter to a king, and in Event 5 the king receives the letter, it seems unlikely that readers expend much effort pondering whether Event 5 depends on Event 1, or on what

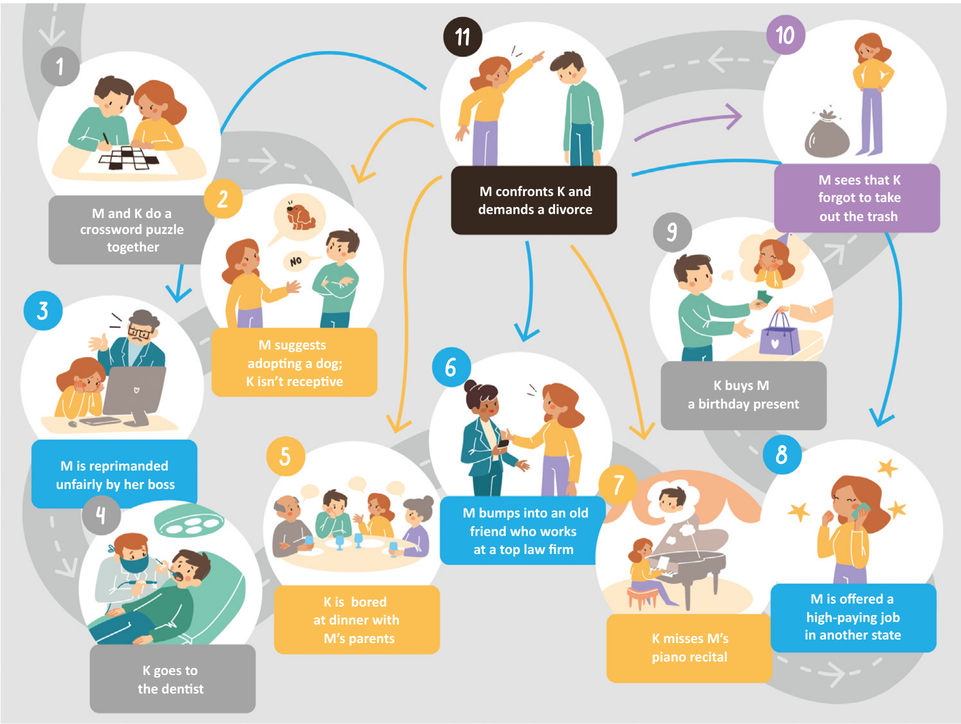

Trends in Cognitive Sciences

Figure 2. Assigning credit to narratively plausible prior events. To which prior events should K assign credit for his current unwanted outcome (#11, black)? If he were to attribute the outcome primarily to the immediately prior event (#10, purple), he would come to the erroneous conclusion that M asked for a divorce because K forgot to take out the trash. Depending solely on temporal contingency to identify relevant states and actions could impede learning of the real causes. By contrast, world knowledge of what types of actions are plausible causes, and familiarity with narrative event sequences associated with romantic relationships, can help identify two sets of possible causes for event #11. Narrative 1 (Neglect; orange): K neglected M's wishes and did not like her family or her hobbies (events #2, #5, and #7); M's building resentment eventually caused her to snap. See Figure 3 in the main text. Narrative 2 (New Opportunity; blue): M was unhappy at her job; a chance meeting with an old friend led to an attractive job offer in another state (events #3, #6, and #8), and M realized that she wanted a fresh start without K. See Figure 3 in the main text. Not all events in K's life contribute to solving the current credit assignment problem. The two narrative sequences (blue and orange) are selected from a larger set that includes events (gray events #1, #4, and #9) that are not plausible causes of event #11. Using narrative structure to narrow and cluster these sets of possible causes can make efficient and effective credit assignment feasible, allowing the people involved to learn how to avoid situations leading to this outcome in the future. (Image credit: Anna Kuzmina, @kiki_arts).

would have happened if Event 1 had never occurred; someone receiving a letter is simply a likely outcome of someone sending a letter. Yet, these two events would be strongly endorsed as being causally linked. Future work is needed to elucidate how causal reasoning during narrative comprehension is related to findings about causal reasoning processes in other domains and under simpler experimental conditions.

Overall, there is ample behavioral evidence that causal connections are an essential determinant of how people understand and remember narratives [36]. This privileged status of causal associations in memory retrieval supports efficient inferences, judgments, and planning [35,37–39]. Indeed, the preference for retrieving – and, by extension, encoding and maintaining – coherent causally linked trajectories may be one way in which individuals can solve otherwise intractable problems of planning and credit assignment in complex, high-dimensional **state spaces** such as are encountered in the real world (Figure 2).

Neural correlates of processing causal structure in narratives

There is limited work directly investigating the effects of causal structure on brain activity when people engage with or produce narratives. However, there are several reasons to suspect that narrative causal structure plays an important role in the activity of brain regions associated with episodic memory, including default mode network (DMN) regions, medial temporal lobe (MTL) cortex, and hippocampus.

The role of mental simulation in narrative comprehension

Causal reasoning might depend on constructing a mental simulation of predicted and counterfactual scenarios, and such simulation is linked to episodic memory systems. The mental model theory of reasoning [40] posits that comprehension of discourse entails constructing a mental simulation of possible scenarios given the currently described event, what has been called a mental microworld [41]. Similarly, Bower and Morrow [42] describe their notion of a mental model or situation model in terms of simulation: 'The bare text is somewhat like a play script that the reader uses like a theater director to construct in imagination a full stage production.' While the quote refers to comprehending text, the same idea applies when the stimulus is a dynamic audiovisual movie: many aspects of the stage production are provided by the film, but many aspects are left to be generated by the audience. For example, if a woman is shown hurriedly packing a suitcase, viewers might construct a mental model that she is running late for a train; if she was previously shown ordering her deadbeat brother to move out, viewers might construct a mental model that she is packing his things. Thus, even an audiovisual movie narrative is still a bare script in some sense, requiring the viewer to do substantial mental work to comprehend the depicted events.

Considerable evidence supports the idea that DMN cortical areas are involved in scene construction and imagery – that is, they may be the site of the stage production in the brain. fMRI studies have shown that imagining future events; for example, visualizing a scene given a text cue such as bus station, elicits activity in medial prefrontal cortex (mPFC), retrosplenial cortex, parahippocampal cortex, and hippocampus [43]. Hippocampal amnesics seem to have difficulty imagining vivid scenes [44,45], and hippocampal activity is elevated during both scene perception and construction [46]. In studies using narrative stimuli, participants who listened to a verbal description of an audiovisual movie they had never seen before exhibited activity patterns in DMN regions that were similar, event by event, to activity patterns in other participants while they actually watched the movie [47]; furthermore, DMN event patterns changed when previously seen movie events needed to be reinterpreted at recollection in light of new information [48]. A complementary set of studies showed that DMN responses exhibit long timescales during narrative comprehension, compatible with slowly changing event information (Box 1). In a study crossing

Box 1. Is there a link between temporal receptive windows and narrative causal coherence?

Several studies have made use of naturalistic stimuli to show that information processing timescales (temporal receptive windows) vary regularly along the cortical hierarchy. In this hierarchy, the shortest timescales are observed in low-level sensory areas, and the longest timescales are observed in higher-order regions more distant from the sensory periphery – overlapping substantially with the DMN. These timescales are often measured using temporal scrambling of naturalistic stimuli (both movies and stories) [81–85]. In DMN cortical areas, responses were affected even when the input was scrambled minimally: randomly re-ordered at the event level (in chunks of 30–60 s). The key result from these studies is that timescales increase systematically across the brain, a principle that cuts across sensory modalities; thus, it was never fully specified what kind of information was being disrupted by the scrambling. In retrospect, a possible explanation is that event-level scrambling specifically affected the integrity of causal structure in the narrative stimuli. In a narrative with randomly reordered events, it may be difficult or impossible to identify causal links between events, leading to impaired comprehension and decreased synchrony across individuals (synchrony being the main measure of interest in these studies). Future work could examine whether the breakage of causal connections in narratives is indeed the key factor that determines whether DMN activity patterns de-synchronize across individuals. Such a finding might suggest that the information accumulation function of long-timescale brain areas subserves the identification of causal connections across temporally extended events.

different proposed features of situation models, posterior DMN regions [posterior medial cortex (PMC), angular gyrus (AG)] and parahippocampal cortex were found to carry information about locations in which events occurred, mPFC coded for event schema, and perirhinal cortex and temporal pole coded for the people, suggesting different regional contributions towards the construction of mental models during narrative comprehension [49]. These studies show how mental models of events could be represented in DMN activity patterns, perhaps populated in part by details retrieved from long-term memory via the MTL and hippocampus.

Neural substrates of association and simulation

Recent work points to a partial dissociation of representations for purely associative and causal structure. Specifically, a carefully designed experiment identified distinct neural correlates of predictions learned according to the Kalman filter and **successor representation** rules, which reflect, respectively, predictive relationships and long-range associative structure, and which were separately associated with activity in middle temporal gyrus and entorhinal cortex [50]. Critically, the middle temporal representation switched depending on the demands of the task, suggesting that individuals may flexibly rely on different horizons of temporal relationships when needed, consistent with the idea that narrative information supports next-step predictions based on multiple scales of continuity. This dissociation is of potentially critical importance for understanding how narratives can support planning. Previous work identified planning-relevant representations in orbitofrontal cortex [51]; the rules that govern transfer between and arbitration among these representations have yet to be conclusively discerned. That each set of regions differentially engages with the DMN may indicate that the DMN supports the transfer of information between representations during offline processing, perhaps supporting the discovery of narrative structure [52,53].

Consistent with this view, several studies point to an overlapping network of regions being important in deliberative decision-making that relies on forward simulation. In particular, a medial prefrontal–hippocampal network has been shown to activate in proportion to the simulation demands of a decision problem [54,55]. In these studies, BOLD signal in this network was observed to increase when the anticipated outcome of a decision was more unpredictable – in other words, when the causal relationship between action and outcome was unclear. Both studies also further showed that such unclear situations yielded reinstatement of potential outcome representations in ventral temporal cortex, consistent with the use of mental simulation based on past experiences. This reinstatement activity also scaled with the difficulty of simulating next-step outcomes, and, in a related study, was shown to correspond to the amount of activation in the hippocampus [56], consistent with the idea that the hippocampus orchestrates goal-directed retrieval of cortical patterns. A related study observed similar activation patterns in correspondence with the simulation demands of intertemporal choices [57], and that this activity (and the resulting decision) could be modified using episodic tags that cued the participant to simulate specific scenarios.

Brain systems that support episodic memory are influenced by narrative causal structure

A few studies have specifically investigated how brain activity relates to causal connections during narrative processing. In one experiment, participants watched temporally scrambled movies and reported moments of subjective comprehension (‘aha’ moments); these coincided with causal connections between events. DMN activity was positively correlated with subjective feelings of comprehension, and exhibited a more functionally integrated and efficient network state at these times [58]. In another study, researchers calculated the degree of causal connectedness (causal centrality) of events in movies based on the network of causal relations between all events [26]. Memory performance and responses in DMN regions (PMC and AG) scaled with causal centrality during spoken recollection of the movies; furthermore, event-offset hippocampal responses were greater for high-centrality events, echoing prior findings that the hippocampus contributes to encoding memories at salient event boundaries [7,59].

Additionally, some neuroimaging studies have made observations relevant to narrative causal structure, even if that was not the explicit focus. In a study investigating hippocampal memory integration processes [60], brain activity was compared across pairs of temporally distant events that were narratively coherent; that is, they could be joined into a two-event story. These were embedded in a longer unrelated narrative. Hippocampal activity patterns were more similar across the narratively coherent event pairs, and reinstated activity one day later predicted the detail of recollection. One interpretation is that the narratively coherent event pairs were causally linked. Other studies have explored the neural correlates of narrative scripts [16,17], thought of as higher-order categories of stories. When participants were presented with a variety of stories occurring in restaurants and airports, the familiar series of causally linked events (e.g., looking at a menu, ordering, eating, and paying) was recovered in DMN regions using a data-driven Hidden Markov Model algorithm during encoding [61] and recall [62]. In a study of how temporal information is tracked during a nonlinear narrative, posterior medial DMN regions were sensitive to temporal jumps, and interevent neural similarity reflected chronological proximity more so than time of presentation [63]. Though between-event causal connections were not explicitly measured, the findings were interpreted as relating to our ability to embed the causal structure of events in a mental model.

Open questions remain as to whether causality-related computations during narrative comprehension and recall take place within the DMN or originate elsewhere, and about the timing of causal reasoning processes during narrative comprehension (Box 2). Nonetheless, narrative causal structure clearly has a substantial impact on neural systems that support episodic memory, and as such it is important to understand this relationship and incorporate it into experimental designs when using narrative stimuli to study the brain.

Box 2. Interpreting extant findings about how narrative causal structure impacts brain activity

DMN activity is correlated with retrieval success and reactivation of episodic memories. In studies using narrative stimuli, is it possible to disentangle whether DMN responses reflect (i) their involvement in causal reasoning processes during encoding and retrieval, as opposed to (ii) enhanced memory encoding and retrieval as a consequence of causality-related computations occurring elsewhere in the brain? One argument against the latter view is that, while DMN activity is positively related to retrieval success, numerous studies have shown that it is lower during encoding for items that are subsequently remembered [86,87]. However, studies showed that DMN activity was greater during encoding when causal recognition [58] or causal connections [26] were higher, and corresponded to higher likelihood of subsequent recall for these events [26]. These observations do not fit neatly with an account of DMN responses to causal processing being merely reflective of enhanced encoding. Meanwhile, studies using non-narrative stimuli have yielded mixed answers to the question of which brain regions support causal reasoning processes, variously implicating networks associated with attention, memory, and executive control (for review, see [88]). To the extent that causal reasoning depends on amodal mental simulations of current and possible events [40] – which, intuitively, may be more true for narrative and daily-life use cases – a key role for the DMN in acts of causal reasoning seems like a plausible proposition.

A related question is whether causal connections in a narrative, when identified retrospectively by expert or naive judgment (see Figure 1 in main text), provide information about the moments in time at which readers or viewers are engaging in causal reasoning processes as they experience the narrative. For instance, many narratives have a prominent chain structure, in which almost every event is judged to have a causal relationship with temporally adjacent events (both immediately prior and immediately following). If every such causal link in a narrative requires causal reasoning, that would imply that these processes are engaged almost continuously during narrative comprehension. Another possible disconnect between retrospective ratings and ongoing causal reasoning is that judgments about events early in a story may be altered in light of information arriving in later parts of the story; this effect could be much greater for certain types of narratives, such as those with surprise endings [48,89] or with intentionally disrupted temporal order [58,63]. Yet another disconnect could arise from individual variability in extracting causal properties of a narrative; for example, in a murder mystery, some viewers might guess the identity of the killer much earlier than others.

A reinforcement learning account of the utility of narratives

Why are narratives treated preferentially in brain and behavior? They may make it easier to perform the computations supporting value-based decision-making by facilitating long-range associations between key actions across events. In the framework of computational reinforcement learning, states are endowed with values proportional to the amount and the likelihood of future rewards that they entail. Neurobiologically, this quantity is thought to be signaled by firing of dopaminergic midbrain neurons **(reward prediction errors; RPEs)**. This value signal is **discounted** in two senses. First, future rewards – those expected to materialize only after additional state transitions – are valued less than immediate rewards, usually according to a constant discount factor applied at each successive step. Second, via the estimated likelihood with which one can expect that reward to follow. A state that has many, equally likely, potential next steps, only a few of which lead to reward, will carry an attenuated value signal from a given future reward when compared to a state with fewer potential successor states. Similarly, rewards obtained only after a long chain of events contribute exponentially less to value signals relayed to the current state. This dispersion of value signals can be mitigated when information is acquired that allows the agent to increase their certainty of a particular next step or eventual outcome. In other words, when they have a higher expectation of a given successor state, the rewards that follow from that successor state will be more heavily credited to the current state. For example, if I walk into a restaurant that I know has one consistently satisfying dish, my expectation of future reward will be higher than when entering a restaurant where the quality varies between experiences.

Narrative structure can impact this second form of discounting. Specifically, experiences that winnow the tree of possible future states should elicit value-like neural signals and behaviors (attentional capture, preferential memory encoding); a prediction that has been borne out in several laboratory experiments. One striking such finding [64] involved monkeys who were offered a choice between two types of cues, each of which led to the same chance of reward at the same delay, but one of which also led to an intermediate cue that also signaled whether or not the trial would ultimately end in reward. Though trials were perfectly matched in every other respect, the monkeys showed a preference for the information-bearing option, and the presentation of this option triggered a response in midbrain dopamine neurons reflecting the increased value – termed an information prediction error. Notably, the prediction error at reward outcome was muted (or enhanced, if the information cue signaled less reward than actually appeared), consistent with a more certain prediction. This is one of several examples where humans and animals express a willingness to pay – in effort, material goods, or even foregone immediate reward – for information about the likelihood of upcoming rewards [65–67], which researchers have argued reflects the **value of information** provided by these cues [68]. Such information may allow individuals to infer causality more quickly.

Narratively structured sequences provide such information at the outset. One can use narrative schema to infer that the causal plausibility (Box 3) of one set of possible action sequences is more likely than another (for instance, stopping to tie my shoes is not likely to have any bearing on the outcome of a meal, whereas the choice of dish should), value signals can spread more efficiently and correctly. This implies greater attention, fluency, and more precise and durable memory for such stimuli (see Outstanding questions for more specific implications). Recent neural and behavioral findings add further support to this view. One fMRI study showed that plausibility alters hippocampal activation during episodic counterfactual thinking: increased plausibility yielded greater hippocampal activation, independent of effort and difficulty [69]. This finding is consistent with preferential evaluation of narratively plausible stimuli. Another study presented individuals with multiple potential causal models to explain a given set of observations, finding that participants appear to show a preference for the more compressed, or parsimonious, option [70], even when such models sacrifice some amount of overall accuracy.

Box 3. How do we learn narrative plausibility?

Narrative plausibility depends on the world knowledge we have accumulated over a lifetime. When we are presented with a scene, the massive space of possible situations that the input (text, video, or the real-world scene before us) could be describing is instantly pruned by world knowledge of what is plausible. In other words, we have a plethora of latent assumptions that seem to guide our interpretations of events:

> Parent describing a book they are reading: 'After the man escaped from prison, he went into hiding for many years.'
> Child: 'What was he hiding behind?'

To an adult, the parent's sentence immediately evokes a class of narratives (a schema) that involve a person concealing their identity, moving to a new city, and tense chase situations with government operatives. To the child, the sentence evokes something like a playground game of hide and seek, albeit a long one, the oddness of which necessitates further questions.

The question of how world knowledge is acquired immediately arises, and indeed one wonders whether the burden of explanation has merely been shifted. It seems unnecessary to defend the existence of world knowledge and common sense. The focus of the current discourse is, rather, on how that world knowledge may be deployed at key moments to guide causal attributions and assign value accurately to past actions or events. Nonetheless, it is a critical open question how narrative plausibility is learned over the course of years, especially early in life.

Recent work shows that by age 3 years, children's behavior is sensitive to narrative structure in fictional stories [90]. Higher-order associations undergo a sensitive period, during which experiencing unpredictable contingencies in the environment may permanently alter mechanisms for interpreting the causal relationships between events [91], with impacts on learning, decision-making, and exploration [92–95]. New, publicly available narrative movie fMRI and electroencephalogram (EEG) datasets with child participants [96,97] also open the door to exploration of developmental questions [98,99].

This result can be explained by the improved certainty of predictions of future events fitting with the given set of causal explanations.

This viewpoint is useful for understanding humans' preference for narratives – especially compact, plausible ones – and potentially for understanding how reinforcement learning is tractable when navigating complex, high-dimensional environments. Recent work [71–73] has shown that the firing of midbrain dopamine neurons, the neural substrate most closely associated with RPE in humans and animals, may in some settings be more precisely described as signaling an update to the contingencies associated with states that precede the current state, implying a predecessor representation akin to the successor described above. In other words, rather than encoding the forward prediction of discounted future reward, the signal conveys the backward prediction of states or features that causally lead to the present outcome – critically, even in the absence of explicit reward [73]. This finding is further supported by work in humans [38] showing that individuals preferentially leverage these backward predictions when a task's state space is diverging: Those tasks in which an agent moving forward in time can occupy more potential states than they have potential starting points. Interestingly, the reverse is also true – individuals are more likely to use forward planning when a task's state space is converging. Both of these findings are consistent with recent theoretical advances that show that it is advantageous for agents to construct compact representations of both backward and forward predictions to support efficient planning, especially in environments where rewards are unstable or state spaces are large but stationary [74].

Taken together, these findings support the idea that behavior and neural reward signals in humans and animals reinforce the acquisition and efficient representation of causal chains of events. We now turn to the question of why such sequences are treated as valuable, in and of themselves.

Credit assignment and the curse of dimensionality

Predictive representations are useful for improving credit assignment. Recent work has shown that efficient credit assignment benefits from selective counterfactual simulation, which is made possible by the sorts of backwards predictive representations described above [37]. However, the current state-of-the-art algorithmic approach to generating such simulations is limited by the **curse of dimensionality** – when state spaces are large, it becomes difficult to prioritize a few out of very many possible preceding (or succeeding) sequences. In this setting, narrative schema could be thought of as a class of prior expectations about the likelihood of a given sequence of events, similar to the set of rich structure priors thought to support inferences about the organization of natural categories [75]. This can enable two critical improvements: First, it allows agents to ignore aspects of the environment that are irrelevant to the likely space of tasks it will encounter – for example, I do not need to expend resources encoding the likelihood that opening my closet door will result in teleportation to a medieval forest. Second, it allows agents to select counterfactual simulations most relevant to the present outcome – when dissatisfied with my lunch, I do not need to simulate what might have happened had I worn different shoes. An agent endowed with priors about causal structure and plausible sequences – in other words, narratives – can use these to increase their certainty about forward and backward contingencies in their environment, improving the efficient spread of value signals that support the selection of plans (Figure 3).

Concluding remarks and future perspectives

Causal structure plays an important role in how we process and remember narratives, both in terms of behavior and brain activity. Consideration of causal variables can shed new light on prior studies and shape future experimentation. Many areas of cognitive research emphasize tasks and stimuli that have no causal structure across trials, and it is worth considering whether this omission leaves a critical gap in our understanding of how the mind and brain process the world. Furthermore, the utility that narratives hold for linking causes to outcomes in complex natural environments, and for making effective plans for the future, may help explain their seemingly privileged status in human memory. These observations point to several avenues for future exploration of how causal structure in narratives, and in daily experience more generally, may influence human learning and health.

Impairments in reasoning are present in certain mental disorders, including schizophrenia and frontotemporal dementia. While schizophrenia is associated with disrupted causal perception

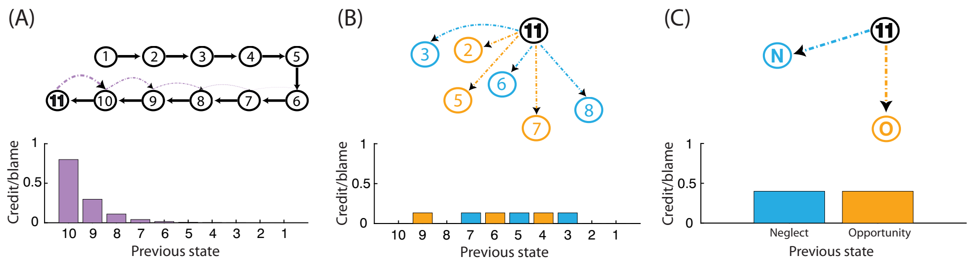

Trends in Cognitive Sciences

Figure 3. The effect of narrative structure on credit assignment. We draw the example from Figure 2 in the main text as a state–space diagram to illustrate how narrative structure can improve the efficacy of credit assignment. (A) The chain of events that preceded the outcome of divorce (#11) is depicted in a strict temporal sequence. Using this representation, an agent might assign credit (blame) mostly to the immediately preceding state (#10), with some decreasing credit assigned to preceding states as well. (B) Using causal plausibility to infer which of the states is responsible for the divorce highlights six states, which in the absence of any further information can be assigned equal credit for the outcome. (C) The problem can be further simplified by realizing that each of the causally plausible events can be understood as one of two narratives that describe the evolution of the relationship. Now, credit can be assigned to each of these two narrative states at a higher level of abstraction, guiding the agent to identify the potential outcomes of future entry into these states.

Outstanding questions

Do naturalistic stimuli with lower overall causal connectivity drive brain responses differently from those with high causal connectivity?

Does the impact of causal structure during narrative processing arise from acts of causal reasoning on the part of the comprehender?

Is causal reasoning during narrative processing different from causal reasoning about physical situations?

Does causal reasoning during narrative processing depend on language (or recruit language processes or neural systems), even for narratives or movies which contain no language?

Is mental imagery an essential aspect of causal reasoning during narrative processing?

Are specific types of causal dependency (temporal priority, operativity, necessity, and sufficiency) important for explaining brain responses during narrative processing?

Can causal structure in narratives inform questions about 'natural' event segmentation grain? That is, should event segmentation be viewed as being driven by the need to identify causal connections across time, which is strategic for later recall?

DMN activity is sensitive to causal variables during narrative processing. Do these responses indicate that the DMN is involved in causal reasoning during encoding and retrieval of narratives, or do they reflect other processes (enhanced memory encoding or retrieval, or more active imagery) that occur as a consequence of causality-related computations occurring elsewhere in the brain?

Regions of the brain with long timescales – at the top of the cortical hierarchy – exhibit responses which are changed when narratives are temporally scrambled at the paragraph or long-event level, as opposed to intact. Are these changes due to disruption of causal structure in the narrative?

Are all narratives equally valued, or only those that fit well-established narrative and reasoning, dementia is more associated with problems in reasoning about psychological and social situations. Both disorders are marked by a decline in episodic memory. Interventions that make use of narratives and storytelling [76,77] may be effective in part because recognizing causal connections is an integral aspect of narrative comprehension; emphasizing causal reasoning in the design of such interventions could increase their benefits.

Another way in which narratives may be valuable to mental health is in improving the adoption of adaptive mental actions, as in psychotherapeutic intervention for treating, for example, rumination or substance use disorder. A major impediment to acquiring adaptive mental actions is that it is difficult to assign credit in mental action space, where many possible and imprecisely defined states can follow a given goal-directed action and precede the desirable outcomes [78]. Narratively structured interventions, for instance, which weave mental actions together with coherent descriptions of a self-identity [79], should make credit assignment more efficient, thus perhaps improving treatment efficacy.

Humans and other animals depend on world knowledge to identify cause–effect relationships between two events or moments in time. While the importance of causal structure for processing and remembering narratives is established, it is unclear how causal reasoning processes are engaged in relation to the network of cause-effect links. Furthermore, distinct cognitive operations and neural substrates could be associated with processing different causal connection types, for example, physical, biological, psychological, social, and structural. These are often present in popular narratives used for modern research, but their taxonomy is not well understood. Consideration of these questions may bridge the gap between ideas about how humans identify and use causal information during complex multi-step experiences, such as narratives or daily life events, with studies and computational models of the neural mechanisms that enable organisms to link actions and outcomes.

schema? How often does a narrative schema need to be experienced before it is incorporated into structural priors? Does the plasticity of these priors change with age and accumulated experience?

Is it the case that simpler narratives are prioritized, as they reduce discounting? In particular, those with fewer intermediate nodes, or which can be chunked into ‘options’ because they contain extended sequences that are consistent with known formats?

What role might prefrontal regions involved in top-down processing play in using causal information to guide cognitive processes during narrative comprehension?

Acknowledgments

The authors gratefully acknowledge the following sources of funding: R21AG072673 (to A.M.B.), R01NS119468 (to A.M.B.; principal investigator: E.R. Chrastil), R01MH133732 (to J.C.). We would like to thank Anna Schapiro for originating ideas, Anna Leshinskaya and Jungsun Yoo for generative discussions, and Miriam Hauptman for insightful comments on an early version of the manuscript.

Citation diversity statement

Recent work in several fields of science has identified a bias in citation practices such that papers from women and other minority scholars are undercited relative to the number of such papers in the field [100–108]. Here, we sought to proactively consider choosing references that reflect the diversity of the field in thought, form of contribution, gender, race, ethnicity, and other factors. First, we obtained the predicted gender of the first and last author of each reference by using databases that store the probability of a first name being carried by a woman [104,109]. By this measure and excluding self-citations to the first and last authors of our current paper, our references contain 14.14% woman(first)/woman(last), 19.1% man/woman, 16.39% woman/man, and 50.36% man/man. This method is limited in that (i) names, pronouns, and social media profiles used to construct the databases may not, in every case, be indicative of gender identity; and (ii) it cannot account for intersex, non-binary, or transgender people. Second, we obtained predicted racial/ethnic category of the first and last author of each reference by databases that store the probability of a first and last name being carried by an author of color [110,111]. By this measure (and excluding self-citations), our references contain 6.66% author of color(first)/author of color(last), 16.86% white author/author of color, 16.89% author of color/white author, and 59.58% white author/white author. This method is limited in that (i) names and Florida Voter Data to make the predictions may not be indicative of racial/ethnic identity; and (ii) it cannot account for Indigenous and mixed-race authors, or those who may face differential biases due to the ambiguous racialization or ethnicization of their names. We look forward to future work that could help us to better understand how to support equitable practices in science.

Declaration of interests

No interests are declared.

References

1. Huk, A. *et al.* (2018) Beyond trial-based paradigms: continuous behavior, ongoing neural activity, and natural stimuli. *J. Neurosci.* 38, 7551–7558
2. Allen, E.J. *et al.* (2022) A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nat. Neurosci.* 25, 116–126
3. Russ, B.E. *et al.* (2023) Temporal continuity shapes visual responses of macaque face patch neurons. *Neuron* 111, 903–914.e3
4. Huth, A.G. *et al.* (2016) Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* 532, 453–458
5. Tang, J. *et al.* (2023) Semantic reconstruction of continuous language from non-invasive brain recordings. *Nat. Neurosci.* 26, 858–866
6. Chen, J. *et al.* (2017) Shared memories reveal shared structure in neural activity across individuals. *Nat. Neurosci.* 20, 115–125
7. Ben-Yakov, A. and Henson, R.N. (2018) The hippocampal film editor: sensitivity and specificity to event boundaries in continuous experience. *J. Neurosci.* 38, 10057–10068
8. Barnett, A.J. *et al.* (2024) Hippocampal-cortical interactions during event boundaries support retention of complex narrative events. *Neuron* 112, 319–330.e7
9. Antony, J.W. *et al.* (2023) Long-term, multi-event surprise correlates with enhanced autobiographical memory. *Nat. Hum. Behav.* 7, 2152–2168
10. Song, H. *et al.* (2021) Neural signatures of attentional engagement during narratives and its consequences for event memory. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2021905118
11. Kim, J. *et al.* (2020) A study in affect: predicting valence from fMRI data. *Neuropsychologia* 143, 107473
12. Chang, L.J. *et al.* (2021) Endogenous variation in ventromedial prefrontal cortex state dynamics during naturalistic viewing reflects affective experience. *Sci. Adv.* 7, eabf7129
13. Ryan, M.-L. (2007) Toward a definition of narrative. In *The Cambridge Companion to Narrative* (Herman, D., ed.), pp. 22–36, Cambridge University Press
14. Radvansky, G.A. and Zacks, J.M. (2011) Event perception. *WIREs. Cogn. Sci.* 2, 608–620
15. Zwaan, R.A. and Radvansky, G.A. (1998) Situation models in language comprehension and memory. *Psychol. Bull.* 123, 162–185
16. Schank, R.C. and Abelson, R.P. (1995) Knowledge and memory: the real story. In *Knowledge and Memory: The Real Story* (Wyer, R.S., ed.), pp. 1–85, Lawrence Erlbaum Associates, Inc
17. Schank, R.C. and Abelson, R.P. (1977) *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*, Lawrence Erlbaum
18. Kintsch, W. and van Dijk, T.A. (1978) Toward a model of text comprehension and production. *Psychol. Rev.* 85, 363–394
19. van Dijk, T. and Kintsch, W. (1983) *Strategies of Discourse Comprehension*, Academic Press, New York
20. Fletcher, C.R. and Chrysler, S.T. (1990) Surface forms, textbases, and situation models: recognition memory for three types of textual information. *Discourse Process.* 13, 175–190
21. Coté, N. *et al.* (1998) Students making sense of informational text: relations between processing and representation. *Discourse Process.* 25, 1–53
22. Mandler, J.M. and Johnson, N.S. (1977) Remembrance of things parsed: story structure and recall. *Cogn. Psychol.* 9, 111–151
23. Trabasso, T. *et al.* (1984) Causal cohesion and story coherence. In *Learning and Comprehension of Text*, pp. 83–111
24. van den Broek, P. (1988) The effects of causal relations and hierarchical position on the importance of story statements. *J. Mem. Lang.* 27, 1–22
25. Trabasso, T. *et al.* (1989) Logical necessity and transitivity of causal relations in stories. *Discourse Process.* 12, 1–25
26. Lee, H. and Chen, J. (2022) Predicting memory from the network structure of naturalistic events. *Nat. Commun.* 13, 4235
27. Bower, G.H. and Clark, M.C. (1969) Narrative stories as mediators for serial learning. *Psychon. Sci.* 14, 181–182
28. Trabasso, T. and Sperry, L.L. (1985) Causal relatedness and importance of story events. *J. Mem. Lang.* 24, 595–611
29. O'Brien, E.J. and Myers, J.L. (1987) The role of causal connections in the retrieval of text. *Mem. Cogn.* 15, 419–427
30. Black, J.B. and Bower, G.H. (1980) Story understanding as problem-solving. *Poetics* 9, 223–250
31. Trabasso, T. and van den Broek, P. (1985) Causal thinking and the representation of narrative events. *J. Mem. Lang.* 24, 612–630
32. Brownstein, A.L. and Read, S.J. (2007) Situation models and memory: the effects of temporal and causal information on recall sequence. *Memory* 15, 730–745
33. Radvansky, G.A. *et al.* (2014) Different kinds of causality in event cognition. *Discourse Process.* 51, 601–618
34. Gerstenberg, T. *et al.* (2017) Eye-tracking causality. *Psychol. Sci.* 28, 1731–1744
35. Lombrozo, T. (2007) Simplicity and probability in causal explanation. *Cogn. Psychol.* 55, 232–257
36. Radvansky, G.A. (2012) Across the event horizon. *Curr. Dir. Psychol. Sci.* 21, 269–272
37. van Hasselt, H. *et al.* (2021) Expected eligibility traces. *Proc. AAAI Conf. Artif. Intell.* 35, 9997–10005
38. Sharp, P.B. and Eldar, E. (2023) Humans adaptively deploy forward and backward planning. *PsyArXiv*, Published online April 11, 2023. https://doi.org/10.31234/osf.io/wdbg4
39. Leshinskaya, A. and Thompson-Schill, S.L. (2019) From the structure of experience to concepts of structure: how the concept "cause" is attributed to objects and events. *J. Exp. Psychol. Gen.* 148, 619–643
40. Johnson-Laird, P.N. (2010) Mental models and human reasoning. *Proc. Natl. Acad. Sci. U. S. A.* 107, 18243–18250
41. Graesser, A.C. *et al.* (1997) Discourse comprehension. *Annu. Rev. Psychol.* 48, 163–189
42. Bower, G.H. and Morrow, D.G. (1990) Mental models in narrative comprehension. *Science* 247, 44–48
43. Addis, D.R. *et al.* (2007) Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia* 45, 1363–1377
44. Hassabis, D. *et al.* (2007) Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl. Acad. Sci. U. S. A.* 104, 1726–1731
45. Mullally, S.L. *et al.* (2012) Scene construction in amnesia: an fMRI study. *J. Neurosci.* 32, 5646–5653
46. Zeidman, P. *et al.* (2014) Constructing, perceiving, and maintaining scenes: hippocampal activity and connectivity. *Cereb. Cortex* 25, 3836–3855
47. Zadbood, A. *et al.* (2017) How we transmit memories to other brains: constructing shared neural representations via communication. *Cereb. Cortex* 27, 4988–5000
48. Zadbood, A. *et al.* (2022) Neural representations of naturalistic events are updated as our understanding of the past changes. *eLife* 11, e79045
49. Reagh, Z.M. and Ranganath, C. (2023) Flexible reuse of cortico-hippocampal representations during encoding and recall of naturalistic events. *Nat. Commun.* 14, 1279
50. Leshinskaya, A. *et al.* (2023) Integration of event experiences to build relational knowledge in the human brain. *Cereb. Cortex* 33, 9997–10012
51. Schuck, N.W. *et al.* (2016) Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* 91, 1402–1412
52. Mildner, J.N. and Tamir, D.I. (2019) Spontaneous thought as an unconstrained memory process. *Trends Neurosci.* 42, 763–777
53. Hahamy, A. *et al.* (2023) The human brain reactivates context-specific past information at event boundaries of naturalistic experiences. *Nat. Neurosci.* 26, 1080–1089
54. Bornstein, A.M. and Daw, N.D. (2013) Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. *PLoS Comput. Biol.* 9, e1003387
55. Doll, B.B. *et al.* (2015) Model-based choices involve prospective neural activity. *Nat. Neurosci.* 18, 767–772
56. Bornstein, A.M. and Norman, K.A. (2017) Reinstated episodic context guides sampling-based decisions for reward. *Nat. Neurosci.* 20, 997–1003

57. Peters, J. and Büchel, C. (2010) Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron* 66, 138–148
58. Song, H. *et al.* (2021) Cognitive and neural state dynamics of narrative comprehension. *J. Neurosci.* 41, 8972–8990
59. Zheng, J. *et al.* (2022) Neurons detect cognitive boundaries to structure episodic memories in humans. *Nat. Neurosci.* 25, 358–368
60. Cohn-Sheehy, B.I. *et al.* (2021) The hippocampus constructs narrative memories across distant events. *Curr. Biol.* 31, 4935–4945.e7
61. Baldassano, C. *et al.* (2018) Representation of real-world event schemas during narrative perception. *J. Neurosci.* 38, 9689–9699
62. Masís-Obando, R. *et al.* (2022) Schema representations in distinct brain networks support narrative memory during encoding and retrieval. *eLife* 11, e70445
63. Grall, C. *et al.* (2023) Neural unscrambling of temporal information during a nonlinear narrative. *Cereb. Cortex* 33, 7001–7014
64. Bromberg-Martin, E.S. and Hikosaka, O. (2009) Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63, 119–126
65. Charpentier, C.J. *et al.* (2018) Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proc. Natl. Acad. Sci. U. S. A.* 115, E7255–E7264
66. Eliaz, K. and Schotter, A. (2007) Experimental testing of intrinsic preferences for nonInstrumental information. *Am. Econ. Rev.* 97, 166–169
67. Iigaya, K. *et al.* (2016) The modulation of savouring by prediction error and its effects on choice. *eLife* 5, e13747
68. Bromberg-Martin, E.S. and Monosov, I.E. (2020) Neural circuitry of information seeking. *Curr. Opin. Behav. Sci.* 35, 62–70
69. Miceli, K. *et al.* (2023) Perceived plausibility modulates hippocampal activity in episodic counterfactual thinking. *Hippocampus* 34, 2–6
70. Kinney, D. and Lombrozo, T. (2023) *Lossy compression and the granularity of causal representation. presented at the NeurIPS 2023 workshop: Information-Theoretic Principles in Cognitive Systems*
71. Garr, E. *et al.* (2023) Mesostriatal dopamine is sensitive to specific cue-reward contingencies. *bioRxiv*, Published online June 7, 2023. https://doi.org/10.1101/2023.06.05.543690
72. Jeong, H. *et al.* (2022) Mesolimbic dopamine release conveys causal associations. *Science* 378, eabq6740
73. Markowitz, J.E. *et al.* (2023) Spontaneous behaviour is structured by reinforcement without explicit reward. *Nature* 614, 108–117
74. Chelu, V. *et al.* (2020) Forethought and hindsight in credit assignment. *arXiv*, Published online October 26, 2020. https://doi.org/10.48550/arXiv.2010.13685
75. Kemp, C. and Tenenbaum, J.B. (2008) The discovery of structural form. *Proc. Natl. Acad. Sci. U. S. A.* 105, 10687–10692
76. Vigliotti, A.A. *et al.* (2019) Evaluating the benefits of the timeslips creative storytelling program for persons with varying degrees of dementia severity. *Am. J. Alzheimers Dis. Other Dement.* 34, 163–170
77. Villar, F. *et al.* (2019) Giving them a voice: challenges to narrative agency in people with dementia. *Geriatrics* 4, 20
78. Hitchcock, P. and Frank, M. (2024) The challenge of learning adaptive mental behavior. *PsyArXiv*, Published online April 15, 2024. https://doi.org/10.31234/osf.io/agprs
79. Pickard, H. (2021) Addiction and the self. *Noûs* 55, 737–761
80. Sasmita, K. and Swallow, K.M. (2023) Measuring event segmentation: an investigation into the stability of event boundary agreement across groups. *Behav. Res.* 55, 428–447
81. Hasson, U. *et al.* (2008) A hierarchy of temporal receptive windows in human cortex. *J. Neurosci.* 28, 2539–2550
82. Lerner, Y. *et al.* (2011) Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *J. Neurosci.* 31, 2906–2915
83. Honey, C.J. *et al.* (2012) Slow cortical dynamics and the accumulation of information over long timescales. *Neuron* 76, 423–434
84. Simony, E. *et al.* (2016) Dynamic reconfiguration of the default mode network during narrative comprehension. *Nat. Commun.* 7, 12141
85. Chien, H.-Y.S. and Honey, C.J. (2020) Constructing and forgetting temporal context in the human cerebral cortex. *Neuron* 106, 675–686.e11
86. Chai, X.J. *et al.* (2014) Development of deactivation of the default-mode network during episodic memory formation. *NeuroImage* 84, 932–938
87. Daselaar, S.M. *et al.* (2004) When less means more: deactivations during encoding that predict subsequent memory. *NeuroImage* 23, 921–927
88. Operskalski, J.T. and Barbey, A.K. (2017) *Cognitive Neuroscience of Causal Reasoning, 1*, Oxford University Press
89. Chen, J. *et al.* (2016) Accessing real-life episodic information from minutes versus hours earlier modulates hippocampal and high-order cortical dynamics. *Cereb. Cortex* 26, 3428–3441
90. Walker, C.M. *et al.* (2015) Learning to learn from stories: children's developing sensitivity to the causal structure of fictional worlds. *Child Dev.* 86, 310–318
91. Harhen, N.C. and Bornstein, A.M. (2024) Interval timing as a computational pathway from early life adversity to affective disorders. *Top. Cogn. Sci.* 16, 92–112
92. Davis, E.P. *et al.* (2022) Early life exposure to unpredictable parental sensory signals shapes cognitive development across three species. *Front. Behav. Neurosci.* 16, 960262
93. Frankenhuis, W.E. and Gopnik, A. (2023) Early adversity and the development of explore–exploit tradeoffs. *Trends Cogn. Sci.* 27, 616–630
94. Spadoni, A.D. *et al.* (2022) Contribution of early-life unpredictability to neuropsychiatric symptom patterns in adulthood. *Depress. Anxiety* 39, 706–717
95. Xu, Y. *et al.* (2023) Childhood unpredictability and the development of exploration. *Proc. Natl. Acad. Sci. U. S. A.* 120, e2303869120
96. Alexander, L.M. *et al.* (2017) An open resource for transdiagnostic research in pediatric mental health and learning disorders. *Sci. Data* 4, 170181
97. Richardson, H. *et al.* (2018) Development of the social brain from age three to twelve years. *Nat. Commun.* 9, 1027
98. Benear, S.L. *et al.* (2023) Setting boundaries: development of neural and behavioral event cognition in early childhood. *Dev. Sci.* 26, e13409
99. Cohen, S.S. *et al.* (2022) Developmental changes in story-evoked responses in the neocortex and hippocampus. *eLife* 11, e69430
100. Mitchell, S.M. *et al.* (2013) Gendered citation patterns in international relations journals. *Int. Stud. Perspect.* 14, 485–492
101. Dion, M.L. *et al.* (2018) Gendered citation patterns across political science and social science methodology fields. *Polit. Anal.* 26, 312–327
102. Caplar, N. *et al.* (2016) Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nat. Astron.* 1, 0141
103. Maliniak, D. *et al.* (2013) The gender citation gap in international relations. *Int. Organ.* 67, 889–922
104. Dworkin, J.D. *et al.* (2020) The extent and drivers of gender imbalance in neuroscience reference lists. *Nat. Neurosci.* 23, 918–926
105. Bertolero, M.A. *et al.* (2020) Racial and ethnic imbalance in neuroscience reference lists and intersections with gender. *bioRxiv*, Published online October 12, 2020. https://doi.org/10.1101/2020.10.12.336230
106. Wang, X. *et al.* (2021) Gendered citation practices in the field of communication. *Ann. Int. Commun. Assoc.* 45, 134–153
107. Chatterjee, P. and Werner, R.M. (2021) Gender disparity in citations in high-impact journal articles. *JAMA Netw. Open* 4, e2114509
108. Fulvio, J.M. *et al.* (2021) Gender (im)balance in citation practices in cognitive neuroscience. *J. Cogn. Neurosci.* 33, 3–7
109. Zhou, D. *et al.* (2020) *Gender Diversity Statement and Code Notebook v1.0. Zenodo*. Published online February 17, 2020. https://doi.org/10.5281/zenodo.3672110
110. Ambekar, A. *et al.* (2009) Name-ethnicity classification from open sources. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 49–58, NY, USA, New York
111. Sood, G. and Laohaprapanon, S. (2018) Predicting race and ethnicity from the sequence of characters in a name. *arXiv*, Published online July 8, 2023. https://doi.org/10.48550/arXiv.1805.02109