



OPEN ACCESS

EDITED BY

Jessecae Marsh,
Lehigh University, United States

REVIEWED BY

Jazmin Cevalco,
University of Buenos Aires, Argentina
Mrinalini Srivastava,
Indian Institute of Technology Delhi, India

*CORRESPONDENCE

Troy M. Houser
✉ thouser@uoregon.edu

RECEIVED 14 April 2025

ACCEPTED 19 September 2025

PUBLISHED 07 October 2025

CITATION

Houser TM (2025) A framework for uniting
space and time in the mind and brain.
Front. Cognit. 4:1611855.
doi: 10.3389/fcogn.2025.1611855

COPYRIGHT

© 2025 Houser. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](#). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

A framework for uniting space and time in the mind and brain

Troy M. Houser*

Department of Psychology, University of Oregon, Eugene, OR, United States

Kant argued that all experience is perceived through the lens of a *a priori* concepts of space and time. That is, Kantian philosophy supposes that knowledge is formatted in terms of space and time. This article argues that space can be reduced to time and thus that the only *a priori* concept used to format knowledge is time. To build this framework, the article focuses on how humans discount time when making intertemporal choices. Celebrated temporal discounting models, such as the exponential and hyperbolic discounting models, are reviewed before arguing in favor of a more ecologically motivated account that suggests that hyperbolic discounting emerges from exponential discounting and uncertainty. The ecological account of temporal discounting is then applied to spatial navigation. Along the way, findings from neurobiology and principles from computational mechanisms are used to substantiate claims. Reducing space to time has important implications in cognitive science and philosophy and can inform a suite of seemingly distinct literatures.

KEYWORDS

temporal discounting, successor representation, exploration-exploitation, cognitive representation, decision-making

Introduction

Experience occurs in both space and time. Thus, everything we do is imbued with a spatiotemporal aspect, yet, as Aristotle remarked many years ago, our mental lives lack an objectively spatial component but seem to abide by the same temporal dynamics as the physical world. Therefore, either space is an innate concept that animals are aware of and understand prior to experience or the idea of space is derived from temporal experience. In this article, I argue for the latter: the view that space is derived from time assumes that space is a latent construct that has to be inferred by the brain, given its lack of direct access to the world. Although the brain does not have direct access to the spatial information in our environment, it has to operate in time, and thus, I assume that the brain does have access to temporal information. This perspective is inspired by philosophical frameworks, such as that espoused by Henri Bergson, Plato's notion of time discussed in *Timaeus*, and Kant's *a priori forms of intuition*, although this article focuses on a neurobiological account.

In what follows, I use research on intertemporal decision-making to elucidate how humans represent time in the mind and brain and how such representations shape decision-making. Specifically, I briefly review a number of computational models that capture intertemporal choices, including the canonical exponential and hyperbolic discounting models, as well as more recently proposed Bayesian accounts of temporal discounting behavior. I then evaluate temporal discounting from an ecological perspective and provide an explanation for how hyperbolic discounting emerges from exponential discounting with heterogeneous discounting rates, an idea originally introduced by Sozou (1998). This idea leads to a plausible neurobiological account of temporal discounting. Assuming that psychological time is represented by a hypothetical neuron population that

discounts time exponentially but with heterogeneity, then this neuron population could also give rise to spatial representations. To explore this possibility, I next review how spatial representations can be derived from the successor representation (SR; Dayan, 1993), which has substantial computational and neurobiological support. Finally, I illustrate the equivalence between a multiscale SR and heterogeneous temporal discounting. I conclude with some potential applications of the proposed framework and testable predictions that follow logically from the guiding principles laid out.

Time

To see how the brain derives its concept of space from its experience with time, I first review how the brain represents time. There are a few prominent theories for how the brain tracks time, including the population and rate coding approaches (MacDonald et al., 2011; Pastalkova et al., 2008). However, here I specifically focus on temporal representation. That is, for example, while a single cell's firing rate may track clock time, it is known that consciously accessible representations of time are anything but perfect. Therefore, I assume that such a cell's firing rate is correlated with, but does not represent, time. Notably, in this article, no claims regarding the objective nature of time or space are made, although, in the discussion, I speculate on how and why the framework proposed is (partially) consistent with Kant's philosophy. For representing time, I turn to cognitive science. The intertemporal choice task, in particular, is an excellent candidate paradigm for studying representations of time.

Background on intertemporal choice

Intertemporal choices are trade-offs between gains and losses occurring at different times (Frederick et al., 2003; Loewenstein and Thaler, 1989). For example, one might have to decide between spending money on a night out that would be immediately rewarding or saving the money to go on an overseas holiday sometime in the future. Decision-making involving outcomes over time is ubiquitous. It characterizes many important life decisions, such as choosing a life partner, saving money for retirement, changing diets for future health, and conserving energy for an upcoming arduous task. Why does everyone not choose health, to save money, and to conserve energy? As was noted in early studies, the only reason for not choosing health, money, and energy is because its payout occurs later. The effect that delaying gratification has on behavior was studied in Hull's (1932) *goal gradients* and the famous Stanford marshmallow experiment (Mischel and Mischel, 1983). These early experiments noted that later rewards are *discounted* (i.e., lessened or reduced in effect). Moreover, they produced considerable research into *how* rewards are discounted with time.

The first formal treatment of temporal discounting emerged from a generalized version of Herrnstein's (1961) matching law, which says that animals allocate their behavior B between two

choices, 1 and 2, in proportion to their rates of reinforcement R :

$$\frac{B_1}{B_2} = \frac{R_1}{R_2}. \quad (1)$$

If we imagine that the behaviors are temporal choices T , then we have the matching law of intertemporal choice:

$$\frac{T_1}{T_2} = \frac{R_1}{R_2}. \quad (2)$$

Baum (1974) generalized the simple matching law to evaluate intertemporal choice:

$$\frac{T_1}{T_2} = b \left(\frac{R_1}{R_2} \right)^s, \quad (3)$$

where b is a bias parameter or a value meant to capture influences beyond reinforcement and s stands for sensitivity to reinforcement. Say one is offered \$10 in 5 days or \$20 in 10 days. If we plug in these numbers, $T_1 = 5$, $T_2 = 10$, $R_1 = 10$, and $R_2 = 20$. This makes the left-hand side of the equation 0.5. When $s = 1$, Equation 3 reduces to Equation 2, and one is simply biased toward one or the other alternative ($\left[\frac{R_1}{R_2} \right]^1 = \left[\frac{R_1}{R_2} \right]$). When $0 < s < 1$, the left-hand side “undermatches” $\frac{R_1}{R_2}$ (e.g., using the previous example, when $s = 0.5$ and $b = 1$, the right-hand side of Equation 3 is 0.71), meaning that not enough behavior is allocated to the immediate reward. When $s > 1$, the left-hand side “overmatches,” meaning not enough behavior is allocated to the later reward. While the matching law of temporal discounting does account for individual sensitivity to reward ratios, it lacked a formulation of subjective value, which is important because experience tells us that preferences are subjective (Peters and Büchel, 2010). Some people are more patient than others (Bornoalova et al., 2009; Daruna and Barnes, 1993; Jauregi et al., 2018), and some people take risks, while others play it safe (Bornoalova et al., 2009; Lejuez et al., 2004; Zuckerman and Kuhlman, 2000). Moreover, the matching law does not assume that animals discount rewards at all, which flies in the face of numerous experiments showing that animals consistently favor immediate gratification even when the reinforcement rates for both options are equal (Rachlin, 2006). Behavioral economic models provided a solution to these problems.

Given that matching is a strategy for reward maximization (Rachlin, 2006; Rachlin et al., 1976), it is also common to view temporal discounting as a utility function. Indeed, Samuelson's (1937) model of discounted utility was the standard model of intertemporal choice for a while. It suggests that future rewards are discounted exponentially with time:

$$V_{LL} = R_{LL} \exp(-kD), \quad (4)$$

where V_{LL} is the subjective value of the larger-later reward, R_{LL} is the larger-later reward offered, D is the time delay until payout, and k is called the temporal discount factor. If $k = 0$, there is no discounting and $V_{LL} = R_{LL}$. As k increases, however, the multiplier that scales R_{LL} , $\exp(-kD)$, shrinks and thereby also shrinks the subjective value of the larger-later reward. For example, assuming $R_{LL} = \$10$ and $D = 5$ days, $V_{LL} = \$9.51$, $\$7.79$, and $\$6.07$ when $k = 0.01$, 0.05 , and 0.10 , respectively. Thus, k is free to vary. The

k parameter became associated with individual levels of patience, or impulse control (Ainslie, 2022), for the smaller k is the more likely one is to wait for the larger-later reward rather than opting for the immediate gratification. Consistent with the centuries-old theory of expected utility, the exponential model of temporal discounting assumes that people's preferences are rational. That is, if one prefers \$100 in a year over \$110 in 2 years, they will also prefer \$100 today over \$110 in a year. This is rational because the delay and reward ratios are the same in both scenarios, meaning that there is nothing new about the second scenario that should change their preference. Of course, Kahneman and Tversky's (1979) explosive prospect theory put an end to belief in humanity's perfect rationality. For example, it is not uncommon for one to prefer \$110 in 2 years over \$100 in a year but also \$100 today over \$110 in a year. These preference reversals can be modeled if rewards are discounted hyperbolically with time. Ainslie (1975) noted that the matching law applied to a single delayed reward would yield a hyperbolic discount curve:

$$V_{LL} \propto \frac{R_{LL}}{D}. \quad (5)$$

The standard formula for hyperbolic discounting became (Mazur, 1987)

$$V_{LL} = \frac{R_{LL}}{1 + kD}. \quad (6)$$

Hyperbolic discount curves decrease more sharply for smaller delays than exponential discount curves (i.e., hyperbolic curves have greater concavity), which was seen as crucial for explaining seemingly irrational addictive behaviors (Ainslie and Monterosso, 2003; Monterosso and Ainslie, 2007). The marked success of exponential and hyperbolic temporal discounting models (Figure 1) at capturing a suite of mental health disorders has led to it being proposed as a transdiagnostic marker of psychopathology (Lempert et al., 2019). Moreover, recent studies have demonstrated the global generalizability of temporal discounting (Rieger et al., 2021; Ruggeri et al., 2022, 2023; Wang et al., 2016). While informative and useful for prediction, deciphering what these computational models mean for our representations of time is not easy.

Ecologically rational intertemporal choice

When faced with two options that do not have an obvious winner, people often deliberate considerably, thinking through possible payoffs and consequences and reflecting on their own goals or histories. This process forces a cognitive agent to construct and search through a mental space of options and outcomes (Bulley and Schacter, 2020). This highly executive process was initially posited to be underpinned by activity in the prefrontal cortex (McClure et al., 2004, 2007) and is contrasted with a separate, impulsive neural substrate in the limbic system, which was suggested to explain the common finding that animals are biased to choose immediate gratification (Ainslie, 2022; McClure et al., 2004, 2007; Peters and Büchel, 2011; Ruggeri et al., 2022). Evidence indeed supports the involvement of the prefrontal

cortex in temporal discounting (Ciaramelli et al., 2021; Kable and Glimcher, 2007; Koban et al., 2023; Peters and Büchel, 2010, 2011). Along similar lines, hippocampal damage in rats shows increased temporal discounting (Cheung and Cardinal, 2005; Mariano et al., 2009; Rawlins, 1985; Rawlins et al., 1985), which is particularly intriguing because hippocampal damage also impairs people's ability to imagine the future (Hassabis et al., 2007). Based on these findings, it has been suggested that the hippocampus contributes to decision-making by projecting oneself into the future (Bar, 2009; Schacter et al., 2007, 2012; Schacter and Addis, 2007). The hippocampus is also known to construct cognitive maps of one's spatial environment using place cells, which are neurons that fire when an animal occupies a specific spatial location (O'Keefe and Dostrovsky, 1971), making the hippocampus a prime candidate for investigating the inference of space from time.

In one of the first studies on the relationship between episodic future thinking and temporal discounting, participants were administered a classical intertemporal choice task involving a series of choices between smaller-sooner and larger-later rewards (Peters and Büchel, 2010). Importantly, the authors introduced a novel condition featuring presentation of episodic cue words referring to real, subject-specific events planned for the day of reward delivery in the intertemporal choice task. It was reasoned that such cues would induce more episodic future thinking and reduce temporal discounting. This is precisely what was found. It was also found that the episodic cue condition induced stronger functional connectivity between the prefrontal cortex and hippocampus (Peters and Büchel, 2010). Thus, it may be the case that the prefrontal cortex draws on mental simulations of the future by the hippocampus to make intertemporal decisions. This neural mechanism is considered evolutionarily advantageous (Boyer, 2008), as it biases choice toward what is better in the long run (Benoit et al., 2011).

Interestingly, this link between temporal discounting and mental simulation has also been used to explain the general bias toward immediate gratification. Given that episodic future thinking decreases temporal discounting and consumption of larger-later rewards is better, we might ask why people do not always employ episodic future thinking strategies to make intertemporal choices. First, episodic future thinking is effortful (Bulley and Schacter, 2020). Thus, to mentally simulate a future reward, one must expend cognitive resources. The more precisely one imagines the future, the more resources that are required (Gershman and Bhui, 2020). As such, a trade-off exists within the trade-off between choices, which is between how much attention to pay toward simulating the future and the magnitude of the future reward. Building on a Bayesian model of hyperbolic discounting (Gabaix and Laibson, 2017), Gershman and Bhui (2020) modeled the magnitude effect of temporal discounting—the finding that people discount time less if the magnitude of the future reward is relatively large—using a model that estimates temporal discounting on a trial-by-trial basis. This model assumes that the mentally simulated future reward is a randomly drawn sample from a Gaussian distribution with mean R_{LL} and variance σ_R^2 , $s \sim N(R_{LL}, \sigma_R^2 D)$. Note that the uncertainty associated with the mental simulation scales with the magnitude of the delay. The key term in this model is the mental simulation noise σ_R^2 , which can be approximated as the variance of all encountered

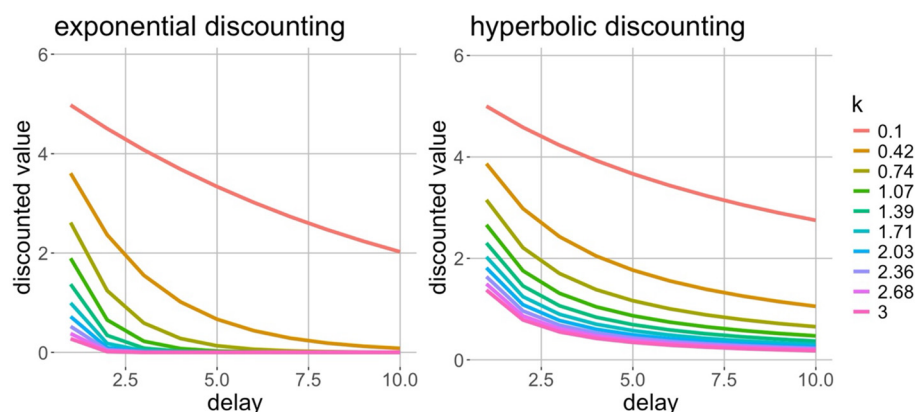


FIGURE 1

Temporal discounting rates. Both plots show a range of temporal delays on the x-axis and hypothetical discounted values of the delayed rewards on the y-axis. Color represents the rate of temporal discounting. On the left, the relationship between discounted value and delay is exponential, and on the right, the relationship is hyperbolic. It can be seen that hyperbolic discounting suggests a more gradual decline in value with delay.

rewards σ_r^2 . This approximation makes implicit that one's prior uncertainty is proportional to the variability of experienced rewards so far. This also captures possible order effects in the intertemporal choice task. Order effects in intertemporal decision-making are rarely considered because k is assumed to be stable across tasks. To capture the magnitude effect (i.e., bring temporal discounting under adaptive control), mental simulation noise σ_R^2 becomes $\frac{\sigma_r^2}{\beta|R|}$, where β , or reward sensitivity, scales the magnitude of the offered rewards. With increasing reward sensitivity, the simulation noise shrinks, which will, in turn, minimize the variability in the sample s . In this model, the discount rate is estimated on each trial as

$$k = \frac{\sigma_R^2}{\sigma_r^2}, \quad (7)$$

which gets plugged into Equation 6, meaning that the discount rate is the ratio of mental simulation noise to observed noise. If mental simulation noise is large compared to observed noise, the discount rate increases. If mental simulation noise is small compared to observed noise, the discount rate decreases, facilitating more patience. Together, temporal discounting is a multifaceted phenomenon shaped by individual differences, as well as reward and uncertainty. However, how more uncertainty biases individuals toward immediate gratification remains unclear. Better yet, *why* does this occur?

One account suggests that premature commitment to suboptimal choices is an adaptive compensatory strategy for high levels of uncertainty that could cause indecision (Salvador et al., 2022). For example, if one has to choose between \$10 now and a draw from a range of possibilities, one might waste precious time trying to reduce the uncertainty of the second option. Consistent with this notion, evolutionary theories have posited the possibility that impulsivity can be adaptive (Stevens and Stephens, 2010). Intertemporal choice is intimately bound up with foraging (Houser, 2024; Namboodiri et al., 2016; Stevens and Stephens, 2010). An animal may wander upon some food, at which point they exploit this resource; however, each bite depletes the resource, and the environment may or may not be changing. At what point does the animal switch to exploring for more

resources? This so-called exploration–exploitation dilemma is an intertemporal choice: exploiting a known resource is analogous to immediate gratification, while exploring incurs a temporal cost but with the hope that a larger reward can be found. Furthermore, exploration is associated with more uncertainty, as is the larger-later reward in intertemporal decision-making. A large reward may exist somewhere in the environment, but it also might not. Stevens and Krebs (1986) showed that when travel times between resources (i.e., larger-later reward delay) are long, exploiting longer travel times is better than short travel times. Environments with plenty of resources will have shorter travel times. Thus, from the perspective of someone in a plentiful environment, overexploitation may appear impulsive, yet overexploitation is adaptive in environments where resources are few and far between (Fenneman and Frankenhuys, 2020). This is fascinating because it means that overexploitation can be a decision-making strategy that depends on thinking about the future, a process not often linked to impulsivity. Together, it is plausible that the so-called present bias in temporal discounting is a function of the risk that the larger-later reward will not be available in the future.

Take, for example, some animal that caches food for the winter. Then, the animal notices a predator lurking in the environment, so the animal digs up its stash and gobbles up the food. On the surface, this behavior might appear myopic; however, it is a rational response to the realization that the predator may steal its stash before winter ends (Houston et al., 1982). If the reward is a future reproductive opportunity, the animal may die before the reproductive act (Iwasa et al., 1984). For humans, an investment opportunity comes with the risk that the stock market collapses. It is thus convenient to define a survival function by specifying the likelihood that a future reward can be realized following some temporal delay, whose inverse can be interpreted as the risk of a reward not being realized. For each unit of time, such risk is known as the hazard rate H :

$$H(D) = \frac{1dS}{SdD}, \quad (8)$$

where S is the survival function. If the hazard rate is constant across all temporal delays, the survival function equals the exponential

discounting model (Sozou, 1998). I already mentioned how animals exhibit preference reversals consistent with hyperbolic discounting, so why should the hazard rate, which is a more ecologically motivated framework, drop off in such a way? Sozou (1998) proposed that animals do assume constant hazard rates, but that uncertainty lies in what the rate is. For simplicity, I assume that the hazard rate can take one of n discrete values, such that the survival function is

$$S(D) = p_1 \exp(-\lambda_1 D) + p_2 \exp(-\lambda_2 D) + \dots + p_n \exp(-\lambda_n D), \quad (9)$$

where p_i is the probability associated with hazard rate λ_i . This is the prior distribution of hazard rates. Generalizing, the Laplace transform of this distribution yields a solution to the case when the hazard rate is continuous:

$$S(D) = \int_0^\infty f(\lambda) \exp(-kD) d\lambda, \quad (10)$$

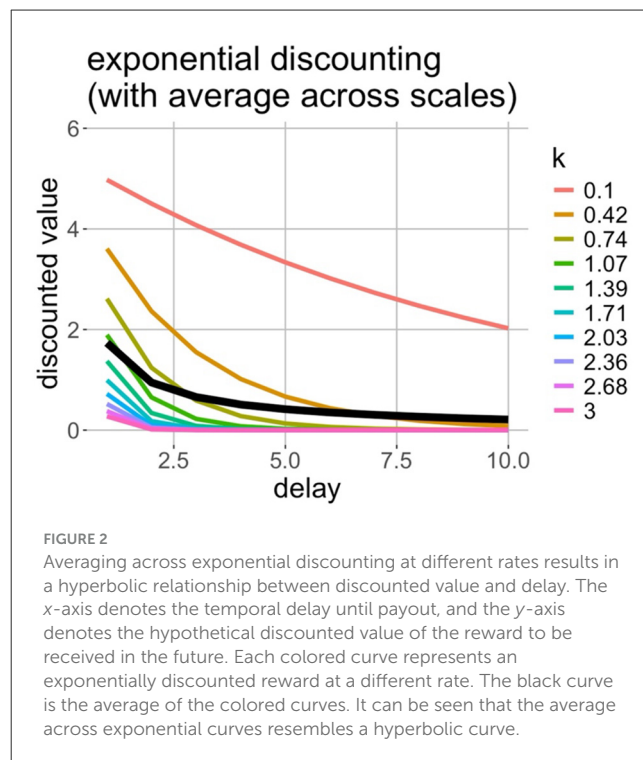
and by setting the prior probability density function $f(\lambda)$ to $\frac{1}{k} \exp(-\frac{\lambda}{k})$,

$$S(D) = \frac{1}{1 + kD}. \quad (11)$$

The remarkable result of this formulation is that taking the Laplace transform of the prior distribution of hazard rates, each one of which individually is consistent with exponential discounting, yields a hyperbolic discounting curve. This is approximately the same as averaging across exponential discount functions, as visualized in Figure 2. In the brain, we might imagine a population of discounting cells, each of which is tuned to a different hazard rate, that collectively produce hyperbolic discounting. Cell types that perform similar functions have been observed in the entorhinal cortex (Bright et al., 2020) and hippocampus (Cao et al., 2021; Howard and Eichenbaum, 2015; Salz et al., 2016). Can this hypothetical population of neurons also give rise to our concept of space?

Space

The hippocampus contains fascinating neurons, called place cells, that fire when an animal occupies a particular spatial location (O'Keefe and Dostrovsky, 1971). Place cells have been hypothesized to serve as the neural substrate for cognitive maps (O'Keefe and Nadel, 1979), and a wealth of empirical evidence supports this notion (Behrens et al., 2018; Bellmund et al., 2018). Experimental work on cognitive maps initially focused on how the brain represents simple environments, such as an open arena (Wilson and McNaughton, 1994). More recently, complex environments resembling many natural environments that are highly structured have been employed (Zhang et al., 2024). Some have suggested that what place cells represent is more akin to a cognitive graph than a cognitive map (Burgess and O'Keefe, 1996; Kumaran and Maguire, 2005; Muller et al., 1996; Peer et al., 2021), the difference being that cognitive graphs feature discrete locations as nodes and distance between nodes equals the number of edges separating the nodes, whereas cognitive maps utilize Euclidean space. The truth is likely



some combination of the two (Peer et al., 2021); however, I believe cognitive graphs make the following discussion more digestible.

If we imagine that space can be represented with nodes and edges (e.g., the Empire State Building and Madison Square Garden can be denoted as two nodes with four edges, or city blocks, between them), then *anything* can be represented spatially. That is, we can just as easily imagine two nodes being red and green separated by three edges (assuming that yellow and orange nodes are in between the red and green nodes). This allows us to talk about color in spatial terms. Indeed, research has elucidated neural underpinnings for odor (Bao et al., 2019), auditory (Aronov et al., 2017), visual (Nau et al., 2018), social (Park et al., 2021), emotion (Qasim et al., 2023), reward (Nitsch et al., 2023), semantic (Viganò et al., 2021), and virtual (Bellmund et al., 2016; Doeller et al., 2010; Julian and Doeller, 2021) spaces, all residing within the entorhinal-hippocampal circuitry. Similarly, conceptual information, which is known to be organized mentally along metric spaces of physical similarity (Jäkel et al., 2008; Kruschke, 1992; Love et al., 2004; Nosofsky, 1986, 1987; Nosofsky and Kruschke, 1992), is also represented in the hippocampus (Bowman et al., 2020; Bowman and Zeithamova, 2018) and entorhinal cortex (Constantinescu et al., 2016). Moreover, navigation in abstract mental spaces recapitulates many characteristics of spatial navigation (Behrens et al., 2018; Bellmund et al., 2018; Giron et al., 2022; Hills et al., 2008, 2013; Meder et al., 2021; Schulz et al., 2018, 2019; Wu et al., 2018, 2020, 2021). The hippocampus representing abstract spaces is also aligned with Eichenbaum's (2017a,b) theory that cognitive maps of one's environment are merely a special case of the hippocampus's much broader role in relational memory (Dusek and Eichenbaum, 1997). Relational memory, however, does not explain

hippocampal involvement in episodic future thought, planning, or intertemporal choice.

An emerging literature attempting to unite the numerous findings regarding cognitive maps has approached the topic from a reinforcement learning perspective. Reinforcement learning is primarily concerned with the problem of reward maximization (Sutton and Barto, 1998). This is a useful lens through which to evaluate the hippocampal function of spatiotemporal representations because reward maximization in the long run requires an ability to both predict the future and represent decision environments (Gershman, 2018). These two requirements of successful reinforcement learning represent the more abstract computational necessities of efficiency and flexibility. Predicting the future allows one to make efficient decisions (e.g., if I know I have a doctor's appointment in the afternoon, I can prioritize small assignments at work because I will not have the time to complete a large assignment). Representing the environment allows one to be flexible (e.g., a cognitive map of the city can facilitate shortcuts or detours if a road closes). A recently revived idea in reinforcement learning, known as the SR (Dayan, 1993), combines both efficiency and flexibility by constructing a predictive map.

The SR learns a predictive map like how the celebrated temporal difference algorithm (Sutton and Barto, 1998) learns a value function. The crucial difference is that SR learning is vector-valued; that is, it iteratively updates values of an entire vector rather than a single point. This vector represents the likelihood of transitioning to any other state, given a starting state. This allows the SR to learn long-term relationships and transition probabilities between states or locations. Specifically, the SR can be constructed via a count-based temporal difference update (Momennejad, 2020):

$$M(s_t) = M(s_t) + \alpha [\mathbb{I}[s_t = s'] + \gamma M(s') - M(s_t)]. \quad (12)$$

Here, M is the SR, $\mathbb{I}[\cdot]$ is an indicator function where every element is 0 except if the argument is true, s_t is the current state, s' is the successor state, α is the learning rate, and γ is the temporal discounting rate. Importantly, M is a matrix with rows and columns equal to the number of states in the environment. Each ij th entry in the matrix represents a transition from state i to state j . Imagine being at home and that you frequent the Empire State Building and Madison Square Garden. M is a 3×3 matrix representing transitions between each of these locations (home, Empire State Building, Madison Square Garden). Say you are home and have plans to see a comedy show at Madison Square Garden. s' represents Madison Square Garden and s_t represents home, making $M(s') - M(s_t) = [0, 0, 0] - [0, 0, 0] = [0, 0, 0]$. When we add the indicator function, we get $[0, 0, 1] + [0, 0, 0] = [0, 0, 1]$. The resulting vector, $[0, 0, 1]$, denotes the fact that you are going to Madison Square Garden. Adding this transition to the original "home vector," $M(s_t)$, yields the vector $[0, 0, 1]$, which in the row of M corresponding to home now means that when at home, you are most likely to go to Madison Square Garden. After numerous transitions, a predictive map is gradually built up that represents the likelihood of all possible transitions between home, Empire State Building, and Madison Square Garden.

By learning transitions among states, the SR can obtain knowledge of space without encoding anything inherently spatial. Say I want to encode the layout of my living room, which is

essentially an open arena. The location one step to my right has to be represented in this cognitive map as one step to my right, the location one step to my left has to be represented as one step to my left, and so on to be able to move about in my living room. Instead of trying to *a priori* figure out what space is, my brain can simply record my transition statistics: I am much more likely to take one than two, three, or four steps to my left, just as I am much more likely to take one than two, three, or four steps to my right. Thus, when in my current location, the SR would infer the likelihood of my next move being one step away. And once I took one step, the next likely move would be to take another single step from that location. The idea of a SR in the mind and brain thus predicts that representing successor states contains implicit knowledge of space without needing to encode anything inherently spatial.

Intriguingly, SRs explain hippocampal place cell activity (Gershman, 2018; Stachenfeld et al., 2017). Place cells have receptive fields that resemble Gaussian tuning curves. That is to say that a place cell is most active in a given location, and its activity decreases as a function of a two-dimensional Gaussian distribution, which is precisely the activation pattern ascribed to place cells by the SR (Figure 3). A two-dimensional Gaussian tuning curve means that a place cell will fire maximally to adjacent locations. Traditional models of place cell activity suggest that this is because place cells are receptive to a patch of space that can be modeled as a two-dimensional Gaussian, whereas the SR suggests that this is because nearby locations are the likeliest successor states. Evidence for predictive maps in the hippocampus has been observed (Ekman et al., 2023; Garvert et al., 2017, 2023; Gershman, 2018; Momennejad, 2020; Sosa and Giocomo, 2021; Stachenfeld et al., 2017).

Predictive maps also lend themselves nicely to the phenomenon of hippocampal replay (Skaggs and McNaughton, 1996), which is when place cells reactivate in the order that they activated previously when foraging at a temporally compressed rate (Foster, 2017; Liu et al., 2019; Wittkuhn et al., 2021). Substantial work shows that hippocampal replay contributes to memory consolidation of salient traveled paths (Deuker et al., 2013; Gagne et al., 2018; Liu et al., 2019; Mattar and Daw, 2018; Momennejad, 2020; Wittkuhn et al., 2021). By encoding multistep predictive maps, the SR can retrace a sequence of transitions that led to reward. This also means that the SR can, in a sense, unroll a sequence of most likely transitions to simulate a future trajectory (Wittkuhn et al., 2021). This latter notion has been termed "preplay" (Buhry et al., 2011; Dragoi and Tonegawa, 2013), and it is involved in planning and prospection (Corneil and Gerstner, 2015; Mattar and Daw, 2018; Ólafsdóttir et al., 2018).

Uniting space and time

In the last section, I neglected to talk about the temporal discounting parameter γ in the SR model. In many behavioral experiments, it is estimated as a free parameter (e.g., Son et al., 2023) to index the amount of generalization over time. How does γ index generalization? To return to the earlier example, which is also visualized in Figure 4, after the show at Madison Square Garden, say you return home. Now s_t is Madison Square Garden and s' is home, making $\mathbb{I}[s_t = s'] + M(s') - M(s_t) =$




FIGURE 3

Example of place cell activity across time. Each grid in the plots can be construed as a matrix of place cell activities, such that each matrix entry represents a single place cell's activation level. Assuming that the location of the animal is at the entry that is brightest, place cell activity decays with distance from this location, resulting in a 2-dimensional multivariate Gaussian distribution (i.e., a radial basis set function). Traditionally, this activation pattern has been interpreted as the tuning of each place cell to space; however, an alternative account is that it represents the likelihood of transitions. That is, one's next step is likely to be to an adjacent spatial location. Thus, place cell activation may denote successor probability.

$[1, 0, 0] + [0, 0, 1] - [0, 0, 0] = [1, 0, 1]$. When we add the resulting vector to the “Madison Square Garden vector,” we get $[1, 0, 1]$. Something interesting happened here. When we first transitioned from home to Madison Square Garden, the SR's home vector became $[0, 0, 1]$, which is intuitive because the 1 marks the likelihood of transitioning to Madison Square Garden. But now after transitioning back home, the Madison Square Garden vector suggests that staying at Madison Square Garden is just as likely as going home, even though you never actually stayed at Madison Square Garden. This is because we never applied γ , which effectively treated γ as equal to 1. γ is bounded between 0 and 1, and the larger it is, the larger its temporal horizon; that is, large γ values update successor states that are multiple steps into the future. In the running example, you went from Madison Square Garden to home. The SR already predicts that, when at home, you will transition to Madison Square Garden. Therefore, because being at Madison Square Garden leads to returning to Madison Square Garden (in two steps), the SR assumes there is a non-zero chance of simply staying at Madison Square Garden. If $\gamma = 0.5$, the Madison Square Garden vector would be $[1, 0, 0.5]$, meaning that multistep predictions are weighted less. This can be interpreted as generalization because two states that end up at the same endpoint will be represented similarly according to the SR. See Figure 5 for how the SR captures cognitive graphs as well.

An option's value is discounted in the SR according to how many steps away it is in the predictive map, according to γ^r , where r is the number of steps, distance, or time points. Because the SR multiplies states by the discount rate at each step, discounting is exponential, and the infinite series can be expressed as

$$M(s_t, s') = \sum_{t=0}^{\infty} \gamma^t \sum_{s_t} p(s' | s_t) \mathbb{I}[s_t = s']. \quad (13)$$

To obtain estimates for the entire transition matrix T , we can express the SR as

$$M = \sum_{t=0}^{\infty} \gamma^t T^t, \quad (14)$$

Or, using the matrix inverse,

$$M = (I - \gamma T)^{-1}, \quad (15)$$

where I is the identity matrix.

The discount rate in the SR can alternatively be expressed as $-\ln(\gamma) = k$, making

$$\gamma^r = \exp(-k \cdot r). \quad (16)$$

Equation 16 is the traditional exponential temporal discounting model using distance instead of temporal delays. Notably, Equation 16 indicates that for any SR discount rate $\gamma \in [0, 1]$, there is an equivalent discount rate k such that the SR's discounting function can be expressed as Samuelson's exponential discounting function. As a result, choosing where to travel to next in one's n -dimensional decision environment can be treated as an intertemporal choice between states in a predictive map, and topological and geometric representations are implicit in the representation of discounted successor states.

Together, it is indeed plausible that we acquire spatial representations from the cognitive processing of time or, more precisely, of sequences, and these predictive maps that can recapitulate physical or abstract spaces are supported by hippocampal activity.

The idea that the same multiscale temporal discounting model can explain predictive maps is mathematically concrete, but does empirical evidence support it? No studies have explicitly tested the connection between spatial and temporal representations using the multiscale temporal discounting model; however, some nascent results hint at its biological plausibility. A particularly telling study tested the extent to which cognitive maps in hippocampal place cell activity can be explained using metrics of hyperbolic vs. Euclidean geometry (Zhang et al., 2023). This study found that hyperbolic geometry fits the data better and that such a representation maximizes the amount of spatial information contained in a population of place cells, making hyperbolic representations evolutionarily advantageous. Moreover, Masuda et al. (2020) found that place cells, in addition to encoding spatial information, encoded temporal delays to larger-later rewards.




FIGURE 4

Successor matrices capture multistep transitions. (A) Traveling from home to Madison Square Garden results in updating the successor matrix. The right side shows a successor matrix after updating. It reflects successor probability (brighter colors = higher likelihood). Specifically, each entry is the probability of transitioning to column j given row i . Thus, because the only transition has been from home to Madison Square Garden, the successor matrix assumes that if the traveler is at home, they will go to Madison Square Garden. (B) Traveling from Madison Square Garden to home results in another update; however, this time, the vector representing Madison Square Garden successors assigns equal likelihood to home and Madison Square Garden. This is because the matrix already assumes that the traveler will transition from home to Madison Square Garden, and thus, being at Madison Square Garden means that they will return to Madison Square Garden in two steps. Thus, the successor matrix captures multistep transitions. (C) This shows the same successor matrix as shown in this (B), but when the temporal discounting rate is set to 0.5, resulting in a reduced update to successor states that are multiple steps into the future.

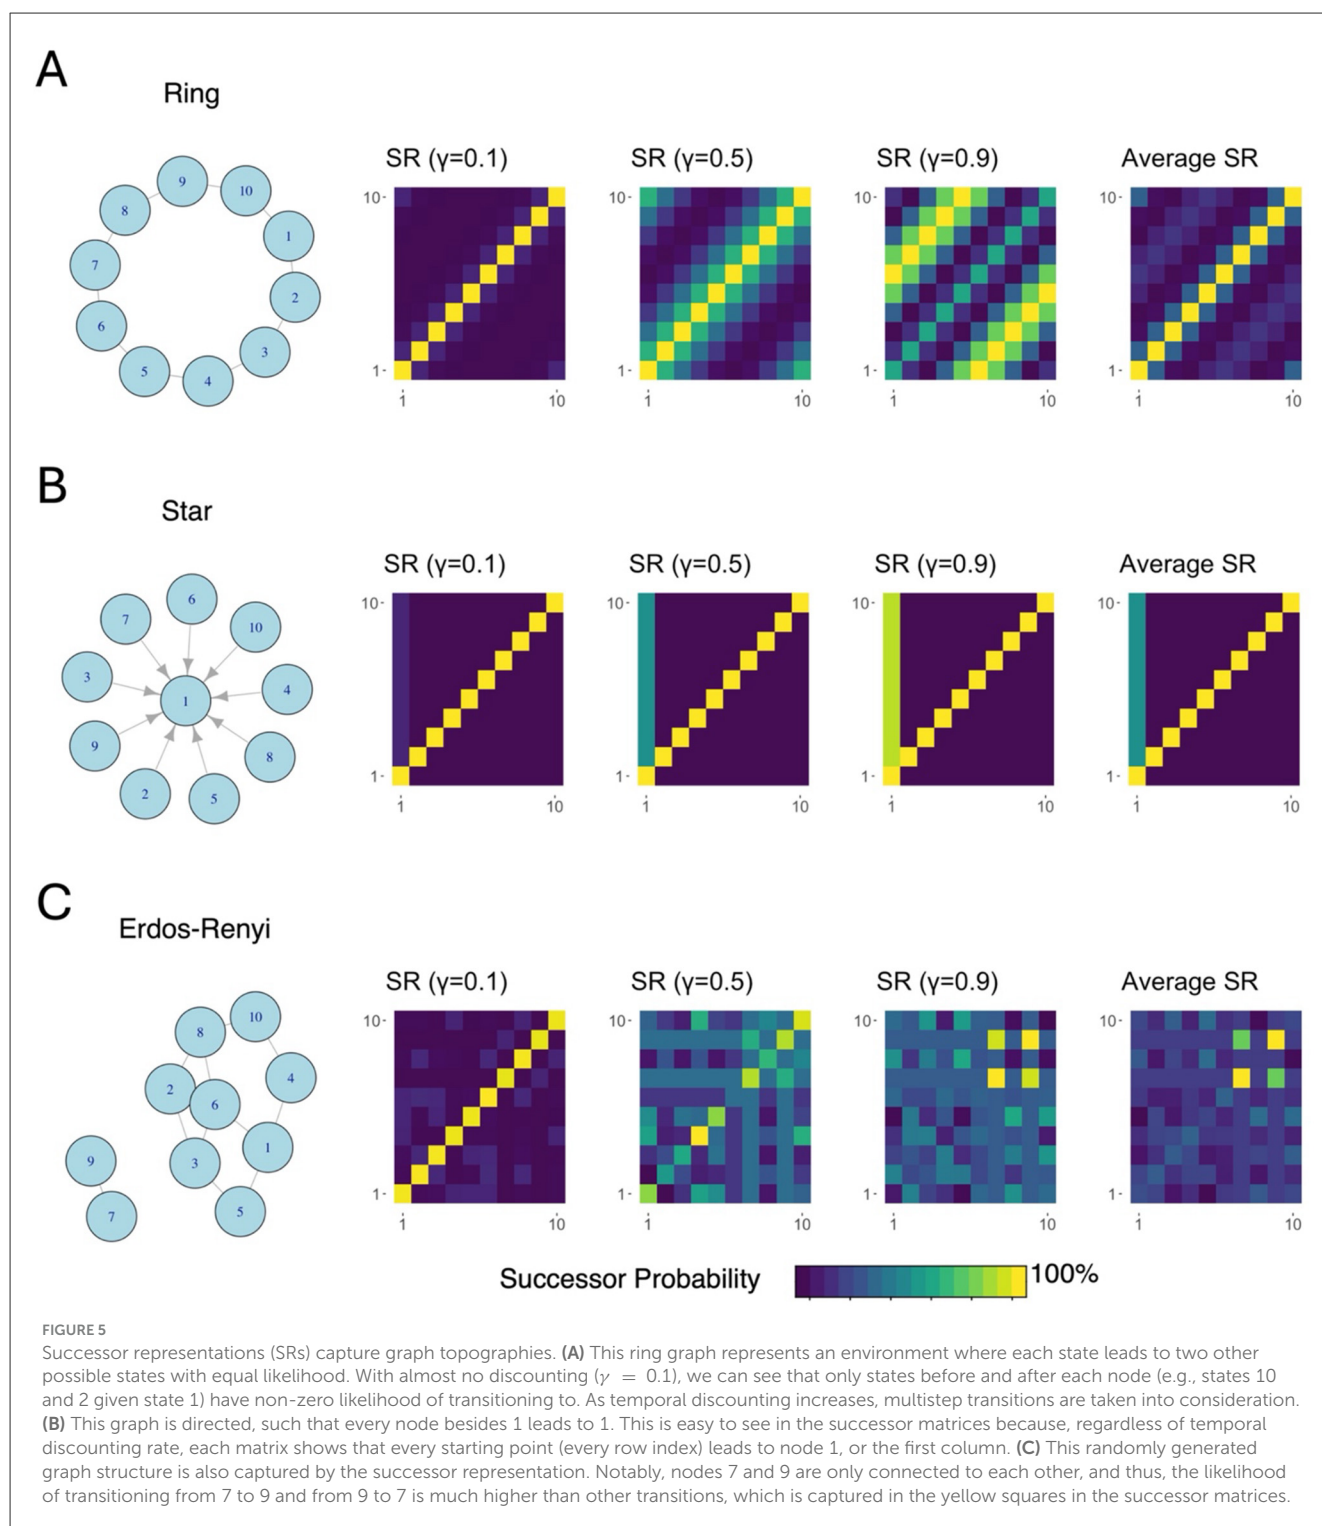
These two studies, in conjunction with the literature on both episodic future thinking and the overlap between foraging and discounting (Harhena and Bornstein, 2023; Houser, 2024; Mikhael and Gershman, 2022; Sozou, 1998; Stevens and Stephens, 2010), entail that it is likely that the hippocampus represents space and time using the same computational principles. Showing that hippocampal spatial representations are inferred from sequences in time will be more challenging to demonstrate, as it requires showing that neural populations transform temporal sequences into predictive maps. Although I remain agnostic as to how this is achieved, a rather simple solution is that each cell in a population has exponential temporal receptive fields with a different decay rate. Thus, each cell would represent the exponential discounting model with a different decay rate. According to Equation 9, the inverse Laplace transform of such a population should yield hyperbolic discounting in space or time. This is precisely what has been predicted and found in the entorhinal cortex (Bright et al., 2020), where so-called temporal context cells track time at distinct rates. It would be interesting to see how many of these temporal context cells are also grid cells to see if the grid code also exhibits hyperbolic geometry.

It is likely that this computational principle of hyperbolic discounting via a population of exponential discounting is utilized in many brain regions due to its efficient nature. It seems specifically useful for representing magnitude information and, thus, may support the analog magnitude system (see Beck, 2015). A graded heterogeneity of time constants has been observed in the cerebellum (Guo et al., 2021) and medial prefrontal cortex (Cao et al., 2024). Further, cells in the retrosplenial cortex have been observed to integrate exponentially discounted past memories, such that proper weighting of individual exponential

discounting rates resulted in a hyperbolic history integration (Danskin et al., 2023).

Discussion

Kant shifted the philosophical tradition of investigating the nature of reality to investigating how knowledge of the nature of reality is possible. A hallmark of Kantian philosophy is that space and time serve as *a priori* intuitions that structure our mental lives and knowledge itself. This article makes the case for reducing these two intuitions to one: time. Consequently, I argued that spatial knowledge is derived from temporal experience. This theoretical framework has wide-ranging applications in cognitive science, particularly for how animals internally represent information. I used the well-researched intertemporal choice paradigm to illustrate how temporal discounting of the future is interchangeable with spatial discounting when considering travel costs associated with exploring one's environment. This overlap was made mathematically concrete by extending Sozou's (1998) temporal discounting model to predictive maps in the hippocampus. Specifically, Sozou (1998) proposed that hyperbolic temporal discounting emerges from an ensemble of exponential discounters because of uncertainty regarding the survival of a reward over some duration. I applied this framework to a multiscale SR (Momennejad and Howard, 2018), which has been observed along the hippocampal longitudinal axis (Brunec and Momennejad, 2022). Together, the theoretical framework outlined here is biologically plausible and supported by evolutionary history. What sort of implications does this framework have?



Insights from the ideas discussed in this article can inform a number of disparate research programs in cognitive science. First, the current framework is consistent with many ideas put forth that attempt to unite all magnitude representations in the mind and brain, such as the Weber–Fechner law (J and Fechner, 1889), a theory of magnitude (Walsh, 2003), and amodal magnitude representations (Nuerk et al., 2005). The idea is that continuous variables are mapped to a generalized

magnitude system in the brain for encoding. While there do appear to be different neural substrates across modalities (Cohen Kadosh et al., 2005; Hamamouche and Cordes, 2019; Sokolowski et al., 2017), evidence supports the notion that each substrate utilizes similar computational principles for representing the information. Specifically, a basis set of exponentially tuned cells with heterogeneous magnitude discounting rates encodes multiscale magnitude representations (Bright et al., 2020; Cao et al.,

2021, 2024; Danskin et al., 2023; Guo et al., 2021; Zhang et al., 2023). It is highly adaptive for the brain's default to be to encode information at multiple scales because it enables flexible behavior. For example, backpacking through Europe likely requires not only a fine-grained map to be able to walk around locally but also a coarse-grained map to board the right trains. The concept of a home requires fine-grained details (e.g., its material, interior design, location) and coarse-grained details (e.g., *a place to sleep, a place to raise children, a place you return to for comfort*). Importantly, an exponential basis set is a universal function approximator (Schaul et al., 2015), meaning it can represent a potentially infinite range of functions, or tuning curves, which could be why animals exhibit considerable diversity in discounting (e.g., hyperbolic, exponential, and quasi-hyperbolic, additive). Notably, one prediction that this article makes is that the functional form of behavior in decision-making paradigms where animals choose between magnitudes depends on uncertainty. When uncertainty is high, the weights for each basis function (assuming decision is a linear combination of weighted basis functions) will be similar, making the decision approximately an average of the basis set, which creates a hyperbolic function. When uncertainty is low, one or a few weights will dominate, and tuning will be exponential. This notion has important testable implications for several fields.

First, assuming lower uncertainty facilitates more exponential discounting, intertemporal choice studies can compare exponential and hyperbolic models when the reward distributions experienced are more or less broad (i.e., uncertain). This notion also suggests that there could be order effects in intertemporal choice tasks, such that the more one learns about the reward distribution, the more their uncertainty decreases, leading to less temporal discounting on later trials. This would be particularly intriguing, given that intertemporal choice tasks are supposed to measure the stable personality trait of impulsivity. Consistent with this notion in general are studies that impose a waiting time between the presentation of the offers and when one must make a choice. Studies found that waiting periods mitigate myopia (McClure et al., 2007), purportedly because one deliberates during this period, which in turn reduces uncertainty (Gilbert and Wilson, 2007).

Together, exploring the notion that time is the most fundamental unit of cognition seems reasonable. For Bergson, every moment of time is gnawed on by the past and impregnated by the future. Every measurement of time or thought placed in spatial formatting is an abstraction from true temporal duration. Being finite and limited in the number of resources animals can employ for representation, animals must efficiently and adaptively approximate this enduring temporality. This article proposes that this is achieved in a population of cells that discounts time exponentially and whose weighted sum can give rise to hyperbolic discounting in decision-making. The computational substrates proposed here were first proposed by Sozou (1998), who elucidated how hyperbolic discounting is rational from an ecological perspective. Sozou's (1998) model was extended to temporal difference (Kurth-Nelson and Redish, 2009) and machine (Fedus et al., 2019) learning. Both models have contributed to understanding value-based decision-making with a population-based account of temporal discounting. Specifically, Kurth-Nelson and Redish (2009) demonstrated that a distributed set of

agents with exponential discounting can approximate hyperbolic discounting. Fedus et al. (2019) extended this model to non-hyperbolic discount functions using a neural network approach. Both models were developed to build on a novel value-based decision-making framework, but their conceptual implications have been relatively unexplored. Moreover, the neurobiological plausibility of these models is unknown. The original contribution of this article is the conceptual unification of spatial representation and temporal discounting. While the framework proposed here does not contradict past models of heterogeneous exponential discounting, it deals with the problem at a complementary level of analysis by outlining the computational task to be solved rather than the algorithmic implementation. Furthermore, the ideas presented in the current article are well grounded in neurobiology.

Intriguingly, this theoretical framework for cognitive representation has broader philosophical implications for integrating Kantian philosophy with cognitive neuroscience. First, to reiterate Bergson's stance, time is a continuous flow, meaning that, while things change, there is never a sharp discretization between events in the real world. Rather, things seamlessly overlap. Our own cells are constantly being replaced by newer cells, but our concept of self is an unbroken stream enduring through time. Consciousness, however, like science, is overwhelmed with this heterogeneity of enduring events and must discretize them to make sense of them. As a solution, consciousness, like science, spatializes time (e.g., the SR). Both artificially segment events into temporal units that can be strung together in a spatial domain like beads on a string, enabling mental (and/or mathematical) manipulation of information. This entails that space is not an *a priori* intuition, as Kant would have it, but rather a construction from *a priori* temporal intuition. Bertrand Russell was a famous empiricist who argued against Kant's idea of *a priori* intuitions (and more broadly his *transcendental idealism*), claiming instead that space and time are objective properties in the world that can be deduced from analysis and logic. While the framework put forth here is largely empirical, it is also meant to encapsulate the spatialization of time and illustrate that even time is represented spatially in the mind, brain, and science. Time itself endures and therefore cannot be known empirically—it can only be lived.

Limitations of this conceptual framework include its computational implementation. While the computational models discussed earlier are robust within their respective fields (i.e., temporal discounting and SR models), how they might be combined in, for example, a neural network architecture to recapitulate neural dynamics is unknown. Furthermore, the philosophical implications have not been thoroughly fleshed out, especially as they relate to the current understanding of the world through the lens of physics. Future work can investigate the extent to which lived time is consistent with what is known in physics.

Data availability statement

The original contributions presented in the study are included in the article/supplementary

material, further inquiries can be directed to the corresponding author.

Author contributions

TH: Conceptualization, Visualization, Writing – original draft, Writing – review & editing.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

- Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychol. Bull.* 82, 463–496. doi: 10.1037/h0076860
- Ainslie, G. (2022). Behavioral construction of the future. *Psychol. Addict. Behav.* 37, 13–24. doi: 10.1037/adb0000853
- Ainslie, G., and Monterosso, J. (2003). “Hyperbolic discounting as a factor in addiction: a critical analysis,” in *Choice, Behavioural Economics and Addiction*, eds. R. E. Vuchinich and N. Heather (Pergamon: Elsevier Science Inc.), 35–69. doi: 10.1016/B978-008044056-9/50043-9
- Aronov, D., Nevers, R., and Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature* 543:7647. doi: 10.1038/nature21692
- Bao, X., Gjorgieva, E., Shanahan, L. K., Howard, J. D., Kahnt, T., and Gottfried, J. A. (2019). Grid-like neural representations support olfactory navigation of a two-dimensional odor space. *Neuron* 102, 1066–1075.e5. doi: 10.1016/j.neuron.2019.03.034
- Bar, M. (2009). The proactive brain: memory for predictions. *Philos. Trans. R. Soc. B* 364:1521. doi: 10.1098/rstb.2008.0310
- Baum, W. M. (1974). On two types of deviation from the matching law: bias and undermatching 1. *J. Exp. Anal. Behav.* 22, 231–242. doi: 10.1901/jeab.1974.22-231
- Beck, J. (2015). Analogue magnitude representations: a philosophical introduction. *Br. J. Philos. Sci.* 66, 829–855. doi: 10.1093/bjps/axu014
- Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., et al. (2018). What is a cognitive map? organizing knowledge for flexible behavior. *Neuron* 100, 490–509. doi: 10.1016/j.neuron.2018.10.002
- Bellmund, J. L. S., Deuker, L., Schröder, T. N., and Doeller, C. F. (2016). Grid-cell representations in mental simulation. *eLife* 5:e17089. doi: 10.7554/eLife.17089.028
- Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., and Doeller, C. F. (2018). Navigating cognition: spatial codes for human thinking. *Science* 362:eaat6766. doi: 10.1126/science.aat6766
- Benoit, R. G., Gilbert, S. J., and Burgess, P. W. (2011). A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *J. Neurosci.* 31, 6771–6779. doi: 10.1523/JNEUROSCI.6559-10.2011
- Bornovalova, M. A., Cashman-Rolls, A., O'Donnell, J. M., Ettinger, K., Richards, J. B., deWit, H., et al. (2009). Risk taking differences on a behavioral task as a function of potential reward/loss magnitude and individual differences in impulsivity and sensation seeking. *Pharmacol. Biochem. Behav.* 93, 258–262. doi: 10.1016/j.pbb.2008.10.023
- Bowman, C. R., Iwashita, T., and Zeithamova, D. (2020). Tracking prototype and exemplar representations in the brain across learning. *eLife* 9, 1–47. doi: 10.7554/eLife.59360
- Bowman, C. R., and Zeithamova, D. (2018). Abstract memory representations in the ventromedial prefrontal cortex and hippocampus support concept generalization. *J. Neurosci.* 38, 2605–2614. doi: 10.1523/JNEUROSCI.2811-17.2018
- Boyer, P. (2008). Evolutionary economics of mental time travel? *Trends Cogn. Sci.* 12, 219–224. doi: 10.1016/j.tics.2008.03.003
- Bright, I. M., Meister, M. L. R., Cruzado, N. A., Tiganj, Z., Buffalo, E. A., and Howard, M. W. (2020). A temporal record of the past with a spectrum of time constants in the monkey entorhinal cortex. *Proc. Natl. Acad. Sci. U.S.A.* 117, 20274–20283. doi: 10.1073/pnas.1917197117
- Brunec, I. K., and Momennejad, I. (2022). Predictive representations in hippocampal and prefrontal hierarchies. *J. Neurosci.* 42, 299–312. doi: 10.1523/JNEUROSCI.1327-21.2021
- Buhry, L., Azizi, A. H., and Cheng, S. (2011). Reactivation, replay, and preplay: how it might all fit together. *Neural Plast.* 2011:203462. doi: 10.1155/2011/203462
- Bulley, A., and Schacter, D. L. (2020). Deliberating trade-offs with the future. *Nat. Hum. Behav.* 4, 238–247. doi: 10.1038/s41562-020-0834-9
- Burgess, N., and O'Keefe, J. (1996). Cognitive graphs, resistive grids, and the hippocampal representation of space. *J. Gen. Physiol.* 107, 659–662. doi: 10.1085/jgp.107.6.659
- Cao, R., Bladon, J. H., Charczyski, S. J., Hasselmo, M. E., and Howard, M. W. (2021). Internally generated time in the rodent hippocampus is logarithmically compressed. *BioRxiv*. doi: 10.1101/2021.10.25.465750
- Cao, R., Bright, I. M., and Howard, M. W. (2024). Ramping cells in rodent mPFC encode time to past and future events via real Laplace transform. *bioRxiv*. doi: 10.1101/2024.02.13.580170
- Cheung, T. H. C., and Cardinal, R. N. (2005). Hippocampal lesions facilitate instrumental learning with delayed reinforcement but induce impulsive choice in rats. *BMC Neurosci.* 6:36. doi: 10.1186/1471-2202-6-36
- Ciamelli, E., De Luca, F., Kwan, D., Mok, J., Bianconi, F., Knyagynska, V., et al. (2021). The role of ventromedial prefrontal cortex in reward valuation and future thinking during intertemporal choice. *eLife* 10:e67387. doi: 10.7554/eLife.67387.sa2
- Cohen Kadosh, R., Henik, A., Rubinsten, O., Mohr, H., Dori, H., Van De Ven, V., et al. (2005). Are numbers special? The comparison systems of the human brain investigated by fMRI. *Neuropsychologia* 43, 1238–1248. doi: 10.1016/j.neuropsychologia.2004.12.017
- Constantinescu, A. O., O'Reilly, J. X., and Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science* 352, 1464–1468. doi: 10.1126/science.aaf0941

Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Cornell, D., and Gerstner, W. (2015). "Attractor network dynamics enable preplay and rapid path planning in maze-like environments," in *Advances in Neural Information Processing Systems, 2015-January* (Massachusetts Institute of Technology Press), 1675–1683.
- Danskin, B. P., Hattori, R., Zhang, Y. E., Babic, Z., Aoi, M., and Komiyama, T. (2023). Exponential history integration with diverse temporal scales in retrosplenial cortex supports hyperbolic behavior. *Sci. Adv.* 9:eadj4897. doi: 10.1126/sciadv. adj4897
- Daruna, J. H., and Barnes, P. A. (1993). "A neurodevelopmental view of impulsivity," in *The Impulsive Client: Theory, Research, and Treatment*, eds. W. G. McCown, J. L. Johnson, and M. B. Shure (American Psychological Association.), 23–37. doi: 10.1037/10500-002
- Dayan, P. (1993). Improving generalization for temporal difference learning: the successor representation. *Neural Comput.* 5, 613–624. doi: 10.1162/neco.1993.5.4.613
- Deuker, L., Olligs, J., Fell, J., Kranz, T. A., Mormann, F., Montag, C., et al. (2013). Memory consolidation by replay of stimulus-specific neural activity. *J. Neurosci.* 33, 19373–19383. doi: 10.1523/JNEUROSCI.0414-13.2013
- Doeller, C. F., Barry, C., and Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature* 463, 657–661. doi: 10.1038/nature08704
- Dragoi, G., and Tonegawa, S. (2013). Distinct preplay of multiple novel spatial experiences in the rat. *Proc. Natl. Acad. Sci. USA* 110, 9100–9105. doi: 10.1073/pnas.1306031110
- Dusek, J. A., and Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proc. Natl. Acad. Sci. USA* 94, 7109–7114. doi: 10.1073/pnas.94.13.7109
- Eichenbaum, H. (2017a). On the integration of space, time, and memory. *Neuron* 95, 1007–1018. doi: 10.1016/j.neuron.2017.06.036
- Eichenbaum, H. (2017b). The role of the hippocampus in navigation is memory. *J. Neurophysiol.* 117, 1785–1796. doi: 10.1152/jn.00005.2017
- Ekman, M., Kusch, S., and de Lange, F. P. (2023). Successor-like representation guides the prediction of future events in human visual cortex and hippocampus. *ELife* 12:e78904. doi: 10.7554/eLife.78904
- Fedus, W., Gelada, C., Bengio, Y., Bellemare, M. G., and Larochelle, H. (2019). *Hyperbolic discounting and learning over multiple horizons*. *arXiv Preprint*. arXiv:1902.06865v3. doi: 10.48550/arXiv.1902.06865
- Fenneman, J., and Frankenhuys, W. E. (2020). Is impulsive behavior adaptive in harsh and unpredictable environments? A formal model. *Evol. Hum. Behav.* 41, 261–273. doi: 10.1016/j.evolhumbehav.2020.02.005
- Foster, D. J. (2017). Replay comes of age. *Annu. Rev. Neurosci.* 40, 581–602. doi: 10.1146/annurev-neuro-072116-031538
- Frederick, S., Loewenstein, G., and O'Donoghue, T. (2003). Time discounting and time preference: a critical review. *J. Econ. Lit.* 40, 351–401.
- Gabaix, X., and Laibson, D. (2017). *Myopia and Discounting*. SSRN (Cambridge, MA: National Bureau of Economic Research). doi: 10.3386/w23254
- Gagne, C., Dayan, P., and Bishop, S. J. (2018). When planning to survive goes wrong: predicting the future and replaying the past in anxiety and PTSD. *Curr. Opin. Behav. Sci.* 24, 89–95. doi: 10.1016/j.cobeha.2018.03.013
- Garvert, M. M., Dolan, R. J., and Behrens, T. E. J. (2017). A map of abstract relational knowledge in the human hippocampal-entorhinal cortex. *ELife* 6:e17086. doi: 10.7554/eLife.17086.021
- Garvert, M. M., Saanum, T., Schulz, E., Schuck, N. W., and Doeller, C. F. (2023). Hippocampal spatio-predictive cognitive maps adaptively guide reward generalization. *Nat. Neurosci.* 26, 615–626. doi: 10.1038/s41593-023-01283-x
- Gershman, S. J. (2018). The successor representation: its computational logic and neural substrates. *J. Neurosci.* 38, 7193–7200. doi: 10.1523/JNEUROSCI.0151-18.2018
- Gershman, S. J., and Bhui, R. (2020). Rationally inattentive intertemporal choice. *Nature Commun.* 11:3365. doi: 10.1038/s41467-020-16852-y
- Gilbert, D. T., and Wilson, T. D. (2007). Projection: experiencing the future. *Science* 317, 1351–1354. doi: 10.1126/science.1144161
- Giron, A. P., Ciranka, S., Schulz, E., Bos, W., Van, D. en, and Wu, C. M. (2022). *Developmental changes in learning resemble stochastic optimization*. *PsyArXiv Preprint*. doi: 10.31234/osf.io/9f4k3
- Guo, C., Huson, V., Macosko, E. Z., and Regehr, W. G. (2021). Graded heterogeneity of metabotropic signaling underlies a continuum of cell-intrinsic temporal responses in unipolar brush cells. *Nature Commun.* 12:5491. doi: 10.1038/s41467-021-22893-8
- Hamamouche, K., and Cordes, S. (2019). Number, time, and space are not singularly represented: Evidence against a common magnitude system beyond early childhood. *Psychon. Bull. Rev.* 26, 833–854. doi: 10.3758/s13423-018-1561-3
- Harhena, N. C., and Bornstein, A. M. (2023). Overharvesting in human patch foraging reflects rational structure learning and adaptive planning. *Proc. Natl. Acad. Sci. U.S.A.* 120:e2216524120. doi: 10.1073/pnas.2216524120
- Hassabis, D., Kumaran, D., and Maguire, E. A. (2007). Using imagination to understand the neural basis of episodic memory. *J. Neurosci.* 27, 14365–14374. doi: 10.1523/JNEUROSCI.4549-07.2007
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement 1, 2. *J. Exp. Anal. Behav.* 4, 267–272. doi: 10.1901/jeab.1961.4.267
- Hills, T. T., Kalf, C., and Wiener, J. M. (2013). Adaptive lévy processes and area-restricted search in human foraging. *PLoS ONE* 8:60488. doi: 10.1371/journal.pone.0060488
- Hills, T. T., Todd, P. M., and Goldstone, R. L. (2008). Search in external and internal spaces. *Psychol. Sci.* 19, 802–808. doi: 10.1111/j.1467-9280.2008.02160.x
- Houser, T. M. (2024). The relationship between temporal discounting and foraging. *Curr. Psychol.* 43, 31149–31158. doi: 10.1007/s12144-024-06716-9
- Houston, A., Kacelnik, A., and McNamara, J. (1982). Some learning rules for acquiring information. *Functional Ontogeny* 1, 140–191.
- Howard, M. W., and Eichenbaum, H. (2015). Time and space in the hippocampus. *Brain Res.* 1621, 345–354. doi: 10.1016/j.brainres.2014.10.069
- Hull, C. L. (1932). The goal-gradient hypothesis and maze learning. *Psychol. Rev.* 39, 25–43. doi: 10.1037/h0072640
- Iwasa, Y., Suzuki, Y., and Matsuda, H. (1984). Theory of oviposition strategy of parasitoids. I. Effect of mortality and limited egg number. *Theor. Popul. Biol.* 26, 205–227. doi: 10.1016/0040-5809(84)90030-3
- J, J., and Fechner, G. T. (1889). Elemente der Psychophysik. *Am. J. Psychol.* 2:669. doi: 10.2307/1411906
- Jäkel, F., Schölkopf, B., and Wichmann, F. A. (2008). Generalization and similarity in exemplar models of categorization: insights from machine learning. *Psychon. Bull. Rev.* 15, 256–271. doi: 10.3758/PBR.15.2.256
- Jauregi, A., Kessler, K., and Hassel, S. (2018). Linking cognitive measures of response inhibition and reward sensitivity to trait impulsivity. *Front. Psychol.* 9:2306. doi: 10.3389/fpsyg.2018.02306
- Julian, J. B., and Doeller, C. F. (2021). Remapping and realignment in the human hippocampal formation predict context-dependent spatial behavior. *Nature Neurosci.* 24, 863–872. doi: 10.1038/s41593-021-00835-3
- Kable, J. W., and Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neurosci.* 10, 1625–1633. doi: 10.1038/nn2007
- Kahneman, D., and Tversky, A. (1979). Prospect theory - an analysis of decision under risk.pdf. *Econometrica* 47, 263–291. doi: 10.2307/1914185
- Koban, L., Lee, S., Schelski, D. S., Simon, M. C., Lerman, C., Weber, B., et al. (2023). An fMRI-based brain marker of individual differences in delay discounting. *J. Neurosci.* 43, 1600–1613. doi: 10.1523/JNEUROSCI.1343-22.2022
- Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychol. Rev.* 99, 22–44. doi: 10.1037/0033-295X.99.1.22
- Kumaran, D., and Maguire, E. A. (2005). The human hippocampus: cognitive maps or relational memory? *J. Neurosci.* 25, 7254–7259. doi: 10.1523/JNEUROSCI.1103-05.2005
- Kurth-Nelson, Z., and Redish, A. D. (2009). Temporal-difference reinforcement learning with distributed representations. *PLoS ONE* 4:af0bf22c83eb. doi: 10.1371/annotation/4a24a185-3eff-454f-9061-af0bf22c83eb
- Lejuez, C. W., Simmons, B. L., Aklin, W. M., Daughters, S. B., and Dvir, S. (2004). Risk-taking propensity and risky sexual behavior of individuals in residential substance use treatment. *Addict. Behav.* 29:35. doi: 10.1016/j.addbeh.2004.02.035
- Lempert, K. M., Steinglass, J. E., Pinto, A., Kable, J. W., and Simpson, H. B. (2019). Can delay discounting deliver on the promise of RDoC? *Psychol. Med.* 49, 190–199. doi: 10.1017/S0033291718001770
- Liu, Y., Dolan, R. J., Kurth-Nelson, Z., and Behrens, T. E. J. (2019). Human replay spontaneously reorganizes experience. *Cell* 178, 640–652.e14. doi: 10.1016/j.cell.2019.06.012
- Loewenstein, G., and Thaler, R. H. (1989). Anomalies: intertemporal choice. *J. Econ. Perspect.* 3, 181–193. doi: 10.1257/jep.3.4.181
- Love, B. C., Medin, D. L., and Gureckis, T. M. (2004). SUSTAIN: a network model of category learning. *Psychol. Rev.* 111, 309–332. doi: 10.1037/0033-295X.111.2.309
- MacDonald, C. J., Lepage, K. Q., Eden, U. T., and Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontinuous events. *Neuron* 71, 737–749. doi: 10.1016/j.neuron.2011.07.012
- Mariano, T. Y., Bannerman, D. M., McHugh, S. B., Preston, T. J., Rudebeck, P. H., Rudebeck, S. R., et al. (2009). Impulsive choice in hippocampal but not orbitofrontal cortex-lesioned rats on a nonspatial decision-making maze task. *Eur. J. Neurosci.* 30, 472–484. doi: 10.1111/j.1460-9568.2009.06837.x
- Masuda, A., Sano, C., Zhang, Q., Goto, H., McHugh, T. J., Fujisawa, S., et al. (2020). The hippocampus encodes delay and value information during delay-discounting decision making. *ELife* 9:e52466. doi: 10.7554/eLife.52466.sa2

- Mattar, M. G., and Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neurosci.* 21, 1609–1617. doi: 10.1038/s41593-018-0232-z
- Mazur, J. E. (1987). “An adjusting procedure for studying delayed reinforcement,” in *Quantitative Analyses of Behavior* (Vol. 5), *The Effect of Delay and of Intervening Events on Reinforcement Value*, eds. M. L. Commons, J. E. Mazur, J. A. Nevin, H. Rachlin (Erlbaum), 55–73.
- McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2007). Time discounting for primary rewards. *J. Neurosci.* 27, 5796–5804. doi: 10.1523/JNEUROSCI.4246-06.2007
- McClure, S. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science* 306:5695. doi: 10.1126/science.1100907
- Meder, B., Wu, C. M., Schulz, E., and Ruggeri, A. (2021). Development of directed and random exploration in children. *Dev. Sci.* 24:e13095. doi: 10.1111/desc.13095
- Mikhael, J. G., and Gershman, S. J. (2022). Impulsivity and risk-seeking as Bayesian inference under dopaminergic control. *Neuropsychopharmacology* 47, 465–476. doi: 10.1038/s41386-021-01125-z
- Mischel, H. N., and Mischel, W. (1983). The development of children's knowledge of self-control strategies. *Child Dev.* 54, 603–619. doi: 10.2307/1130047
- Momennejad, I. (2020). Learning structures: predictive representations, replay, and generalization. *Curr. Opin. Behav. Sci.* 32, 155–166. doi: 10.1016/j.cobeha.2020.02.017
- Momennejad, I., and Howard, M. (2018). Predicting the future with multi-scale successor representations. *BioRxiv*. doi: 10.1101/449470
- Monterosso, J., and Ainslie, G. (2007). The behavioral economics of will in recovery from addiction. *Drug Alcohol Depend.* 90, S100–S111. doi: 10.1016/j.drugalcdep.2006.09.004
- Muller, R. U., Stead, M., and Pach, J. (1996). The hippocampus as a cognitive graph. *J. Gen. Physiol.* 107, 663–694. doi: 10.1085/jgp.107.6.663
- Namoodiri, V. M. K., Levy, J. M., Mihalas, S., Sims, D. W., and Shuler, M. G. H. (2016). Rationalizing spatial exploration patterns of wild animals and humans through a temporal discounting framework. *Proc. Natl. Acad. Sci. U.S.A.* 113, 8747–8752. doi: 10.1073/pnas.1601664113
- Nau, M., Navarro Schröder, T., Bellmund, J. L. S., and Doeller, C. F. (2018). Hexadirectional coding of visual space in human entorhinal cortex. *Nature Neurosci.* 21, 188–190. doi: 10.1038/s41593-017-0050-8
- Nitsch, A., Garvert, M. M., Bellmund, J. L. S., Schuck, N. W., and Dollar, C. F. (2023). Grid-like entorhinal representation of an abstract value space during prospective decision making. *bioRxiv*. doi: 10.1101/2023.08.02.548378
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *J. Exp. Psychol. Gen.* 115, 39–61. doi: 10.1037//0096-3445.115.1.39
- Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *J. Exp. Psychol. Learn. Mem. Cogn.* 13, 87–108. doi: 10.1037/0278-7393.13.1.87
- Nosofsky, R. M., and Kruschke, J. K. (1992). Investigations of an exemplar-based connectionist model of category learning. *Psychol. Learn. Motiv.* 28, 207–250. doi: 10.1016/S0079-7421(08)60491-0
- Nuerk, H. C., Wood, G., and Willmes, K. (2005). The universal SNARC effect: The association between number magnitude and space is amodal. *Exp. Psychol.* 52:187. doi: 10.1027/1618-3169.52.3.187
- O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34, 171–175. doi: 10.1016/0006-8993(71)90358-1
- O'Keefe, J., and Nadel, L. (1979). The hippocampus as a cognitive map. *Behav. Brain Sci.* 2, 487–494. doi: 10.1017/S0140525X00063949
- Ólafsdóttir, H. F., Bush, D., and Barry, C. (2018). The role of hippocampal replay in memory and planning. *Curr. Biol.* 28, R37–R50. doi: 10.1016/j.cub.2017.10.073
- Park, S. A., Miller, D. S., and Boorman, E. D. (2021). Inferences on a multidimensional social hierarchy use a grid-like code. *Nat. Neurosci.* 24, 1292–1301. doi: 10.1038/s41593-021-00916-3
- Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science* 321:5894. doi: 10.1126/science.1159775
- Peer, M., Brunec, I. K., Newcombe, N. S., and Epstein, R. A. (2021). Structuring knowledge with cognitive maps and cognitive graphs. *Trends Cogn. Sci.* 25, 37–54. doi: 10.1016/j.tics.2020.10.004
- Peters, J., and Büchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-midtemporal interactions. *Neuron* 66, 138–148. doi: 10.1016/j.neuron.2010.03.026
- Peters, J., and Büchel, C. (2011). The neural mechanisms of inter-temporal decision-making: Understanding variability. *Trends Cogn. Sci.* 15, 227–239. doi: 10.1016/j.tics.2011.03.002
- Qasim, S. E., Reinacher, P. C., Brandt, A., Schulze-Bonhage, A., Kunz, L., and Qasim, S. (2023). Neurons in the human entorhinal cortex map abstract emotion space. *bioRxiv*. doi: 10.1101/2023.08.10.552884
- Rachlin, H. (2006). Notes on discounting. *J. Exp. Anal. Behav.* 85, 231–242. doi: 10.1901/jeab.2006.85-05
- Rachlin, H., Green, L., Kagel, J. H., and Battalio, R. C. (1976). Economic demand theory and psychological studies of choice. *Psychol. Learn. Motiv.* 10, 129–154. doi: 10.1016/S0079-7421(08)60466-1
- Rawlins, J. N. P. (1985). Associations across time: the hippocampus as a temporary memory store. *Behav. Brain Sci.* 8, 479–497. doi: 10.1017/S0140525X00001291
- Rawlins, J. N. P., Feldon, J., Ursin, H., and Gray, J. A. (1985). Resistance to extinction after schedules of partial delay or partial reinforcement in rats with hippocampal lesions. *Exp. Brain Res.* 59, 273–281. doi: 10.1007/BF00230907
- Rieger, M. O., Wang, M., and Hens, T. (2021). Universal time preference. *PLoS ONE* 16:245692. doi: 10.1371/journal.pone.0245692
- Ruggeri, K., Ashcroft-Jones, S., Abate Romero Landini, G., Al-Zahli, N., Alexander, N., Andersen, M. H., et al. (2023). The persistence of cognitive biases in financial decisions across economic groups. *Sci. Rep.* 13:10329. doi: 10.1038/s41598-023-36339-2
- Ruggeri, K., Panin, A., Vdovic, M., Večkalov, B., Abdul-Salaam, N., Achterberg, J., et al. (2022). The globalizability of temporal discounting. *Nat. Hum. Behav.* 6, 1386–1397. doi: 10.1038/s41562-022-01392-w
- Salvador, A., Arnal, L. H., Vinckier, F., Domenech, P., Gaillard, R., and Wyart, V. (2022). Premature commitment to uncertain decisions during human NMDA receptor hypofunction. *Nature Commun.* 13:338. doi: 10.1038/s41467-021-27876-3
- Salz, D. M., Tiganj, Z., Khasnabish, S., Kohley, A., Sheehan, D., Howard, M. W., et al. (2016). Time cells in hippocampal area CA3. *J. Neurosci.* 36, 7476–7484. doi: 10.1523/JNEUROSCI.0087-16.2016
- Samuelson, P. A. (1937). A note on measurement of utility. *Rev. Econ. Stud.* 4, 155–161. doi: 10.2307/2967612
- Schacter, D. L., and Addis, D. R. (2007). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philos. Trans. R Soc. Lond. B Biol. Sci.* 362, 773–786. doi: 10.1098/rstb.2007.2087
- Schacter, D. L., Addis, D. R., and Buckner, R. L. (2007). Remembering the past to imagine the future: the prospective brain. *Nat. Rev. Neurosci.* 8, 657–661. doi: 10.1038/nrn2213
- Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., and Szpunar, K. K. (2012). The future of memory: remembering, imagining, and the brain. *Neuron* 76, 677–694. doi: 10.1016/j.neuron.2012.11.001
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. (2015). Universal value function approximators. *Proc. Mach. Learn. Res.* 37, 1312–1320. Available online at: <https://proceedings.mlr.press/v37/schaul15.html>
- Schulz, E., Wu, C. M., Huys, Q. J. M., Krause, A., and Speekenbrink, M. (2018). Generalization and search in risky environments. *Cogn. Sci.* 42, 2592–2620. doi: 10.1111/cogs.12695
- Schulz, E., Wu, C. M., Ruggeri, A., and Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychol. Sci.* 30, 1561–1572. doi: 10.1177/0956797619863663
- Skaggs, W. E., and McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271, 1870–1873. doi: 10.1126/science.271.5257.1870
- Sokolowski, H. M., Fias, W., Bosah Ononye, C., and Ansari, D. (2017). Are numbers grounded in a general magnitude processing system? A functional neuroimaging meta-analysis. *Neuropsychologia* 105, 50–69. doi: 10.1016/j.neuropsychologia.2017.01.019
- Son, J. Y., Bhandari, A., and FeldmanHall, O. (2023). Abstract cognitive maps of social network structure aid adaptive inference. *Proc. Natl. Acad. Sci. U.S.A.* 120:e2310801120. doi: 10.1073/pnas.2310801120
- Sosa, M., and Giocomo, L. M. (2021). Navigating for reward. *Nat. Rev. Neurosci.* 22, 472–487. doi: 10.1038/s41583-021-00479-z
- Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proc. R. Soc. B* 265, 2015–2020. doi: 10.1098/rspb.1998.0534
- Stachenfeld, K. L., Botvinick, M. M., and Gershman, S. J. (2017). The hippocampus as a predictive map. *Nat. Neurosci.* 20, 1643–1653. doi: 10.1038/nn.4650
- Stevens, J. R., and Stephens, D. W. (2010). “The adaptive nature of impulsivity,” in *Impulsivity: The Behavioral and Neurological Science of Discounting*, eds. G. J. Madden and W. K. Bickel (American Psychological Association), 361–387. doi: 10.1037/12069-013

- Stevens, T. A., and Krebs, J. R. (1986). Retrieval of stored seeds by Marsh Tits *Parus palustris* in the field. *Ibis* 128, 513–525. doi: 10.1111/j.1474-919X.1986.tb02703.x
- Sutton, R. S., and Barto, A. G. (1998). Reinforcement learning: an introduction. *IEEE Trans. Neural Netw.* 9, 1054–1054. doi: 10.1109/TNN.1998.712192
- Viganò, S., Rubino, V., Soccio, A. Di, Buiatti, M., and Piazza, M. (2021). Grid-like and distance codes for representing word meaning in the human brain. *NeuroImage* 232:117876. doi: 10.1016/j.neuroimage.2021.117876
- Walsh, V. (2003). A theory of magnitude: common cortical metrics of time, space and quantity. *Trends Cogn. Sci.* 7, 483–488. doi: 10.1016/j.tics.2003.09.002
- Wang, M., Rieger, M. O., and Hens, T. (2016). How time preferences differ: evidence from 53 countries. *J. Econ. Psychol.* 52, 115–135. doi: 10.1016/j.joep.2015.12.001
- Wilson, M. A., and McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science* 265, 676–679. doi: 10.1126/science.8036517
- Wittkuhn, L., Chien, S., Hall-McMaster, S., and Schuck, N. W. (2021). Replay in minds and machines. *Neurosci. Biobehav. Rev.* 129, 367–388. doi: 10.1016/j.neubiorev.2021.08.002
- Wu, C. M., Schulz, E., Garvert, M. M., Meder, B., and Schuck, N. W. (2020). Similarities and differences in spatial and nonspatial cognitive maps. *PLoS Comput. Biol.* 16:e1008149. doi: 10.1371/journal.pcbi.1008149
- Wu, C. M., Schulz, E., and Gershman, S. J. (2021). Inference and search on graph-structured spaces. *Comput. Brain Behav.* 4, 125–147. doi: 10.1007/s42113-020-00091-x
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., and Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nat. Hum. Behav.* 2, 915–924. doi: 10.1038/s41562-018-0467-4
- Zhang, H., Rich, P. D., Lee, A. K., and Sharpee, T. O. (2023). Hippocampal spatial representations exhibit a hyperbolic geometry that expands with experience. *Nat. Neurosci.* 26, 131–139. doi: 10.1038/s41593-022-01212-4
- Zhang, T., Rosenberg, M., Jing, Z., Perona, P., and Meister, M. (2024). Endotaxis: a neuromorphic algorithm for mapping, goal-learning, navigation, and patrolling. *ELife* 12:RP84141. doi: 10.7554/eLife.84141.3
- Zuckerman, M., and Kuhlman, D. M. (2000). Personality and risk-taking: common biosocial factors. *J. Personality* 68, 999–1029. doi: 10.1111/1467-6494.00124