

The neural architecture of prediction over a continuum of spatiotemporal scales

Michael T Gravina and Per B Sederberg



Theories of time and space in memory have traditionally focused on their role in dividing experience into discrete episodes, despite the arbitrary nature of these divisions. We offer an alternative characterization that focuses on the fundamentally *predictive* role of perception and memory. In this account, perceptual hierarchies in sensory cortex detect patterns of feature-change across a logarithmic continuum of scales in time and space, which allows them to efficiently converge on nuanced, yet short-range, predictions of the present situation. Time and space emerge from this continuum as representations of feature-distance that provide a measure of the relevance of non-simultaneous experiences, allowing for long-range associations, mental time-travel, and predictions that go far beyond the immediate moment. This reframing of the nature and role of time and space in memory has implications for both the interpretation of existing findings and the design of future experiments.

Address

Department of Psychology, The Ohio State University, USA

Corresponding author: Sederberg, Per B (seederberg.1@osu.edu)

Current Opinion in Behavioral Sciences 2017, **17**:194–202

This review comes from a themed issue on **Memory in time and space**

Edited by **Lila Davachi** and **Neil Burgess**

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 3rd October 2017

<http://dx.doi.org/10.1016/j.cobeha.2017.09.001>

2352-1546/© 2017 Elsevier Ltd. All rights reserved.

Introduction

Humans are drawn to strict divisions and clear categories, which help us to simplify the otherwise intractable complexities we encounter in the world. Yet these same simplifying constructs may become a stumbling block to true comprehension when they impoverish *meaningful* complexity. A foundational concept in the field of memory is the categorical distinction between episodic and semantic memory [1]. However, the very notion of an episodic memory presupposes that an experience is bound to a specific time and space to produce discrete ‘episodes’ [2]. This distinction rests on deeper assumptions that ‘time’ and ‘space’ actually *exist* in the brain as

representational scaffolds onto which experience can be bound. Moreover, the temporal and spatial components of episodic memory are frequently treated in isolation [3], even though it has not been firmly established that they are meaningfully separable.

We argue here that the apparent division of experience into discrete episodes is actually an over-simplification within the basic mechanisms of perception and memory. The flow of experience is continuous across temporal and spatial scales, from milliseconds to decades, and from millimeters to hundreds of kilometers. Given the need to adapt predictively across this full continuum of scales, the basic representations of our experience, and our episodic memories for that experience, must also span this continuum. This proposition does not imply that there are no meaningful boundaries in our experience (see [4] for discussion), but only that the establishment of a spectrum of spatiotemporal scales necessarily precedes, and provides the fundamental substrate for, the definition and identification of such boundaries.

In discussing this idea we also challenge the current discourse on representation in hierarchical sensory cortical streams, such as the ventral temporal stream, in order to move away from models of discrete regional specialization toward a continuous spectrum of scale-sensitivities. Finally, we position the neural representations of time and space as emergent, rather than elemental, properties in the brain, founded on a gradient of experiential scales established in the architecture of the medial temporal lobe.

To ground this proposal, we must first discuss how our senses and all layers of perceptual processing are basically change-detectors operating on continuous streams of low-level features. From this starting point, we must solve the evolutionary challenge of preparing adaptively for the future by making predictions on these feature-changes. Detailing how we accomplish this *efficiently* will lead us to the alternate memory paradigms that we champion here.

Prediction machines

The function of perception is to predict changes in sensory streams

We are accustomed to think of the brain primarily as a device for recognizing and responding to the higher-level wholes of exogenous experience: not only object recognition [5], but also recognition of scenes, events, and the subjectively ‘real’ dimensions of time and space. In this

framework, sensory processing streams in cortex are conceptualized as hierarchies of distinct processing centers, each of which responds to a discrete higher-order category. For example, observations that anatomical loci in the ventral temporal cortex respond to specific visual categories (e.g., lateral occipital cortex, LOC, for objects [6], parahippocampal place area, PPA, for scenes [7], among others) have been taken as evidence that the primary role of these loci is to *represent* those categories [8].

But a living organism does not have direct access to any of these complex structures in its environment. It must extract them from experience mediated entirely by the streams it receives through its sensory receptors, each of which conveys only the intensity of a simple sensory *feature*, like the luminance of light or the pressure of dermal contact, as it varies continuously over time. Thus, at the input level, neither category, nor time, nor space, nor any of the other high-level ‘bins’ of content exist to the brain, only the content itself (a set of continuous, single-feature inputs). Furthermore, as we argue below, recognizing high-level phenomena is not even perception’s fundamental *objective*, but is rather an instrumental byproduct of a living system’s attempt to predict relevant fluctuations in these sensory streams in an energy-efficient manner.

The evolutionary fitness of a living system depends on its ability to efficiently reduce its entropy through interactions with its environment. Even though the universe exhibits a global trend toward increasing entropy, a living system can maintain or even reduce its own entropy, and therefore persist and propagate in the face of constant, disordering external perturbations, by identifying and exploiting *likely* environmental changes. In other words, they can proactively adapt themselves to survive energetic fluctuations from the environment, thereby *correlating* their own behavior with the environment [9–11]. Thus, the environment can be thought of as a state machine, and the organism’s principal adaptive challenge is to infer and encode that state machine’s transition probability structure, using its set of sensory feature inputs as the only available proxy for learning those states and transitions [12]. The resulting correlation between information stored in the system and probabilistically likely regularities in the environment can be considered the rudiment of memory that is harnessed by a living system to guide behavior.

So the primary function of perception is not actually to represent familiar categories of *content*, but rather to register and predict *changes* in that content. This accords well with the established principle that repetitive content is disregarded through neural habituation, with the retained information representing only the change from what was predicted [13]. Thus, every aspect of the

architecture of our sensory processing must ultimately serve to *predict the future of a continuous stream of sensory features based on changes it has observed in that stream in the past*.


Perceptual cortex differentiates scales, not categories

We therefore have reason to doubt the characterization of perceptual hierarchies as chains of loci representing discrete categories and properties. To the extent perceptual hierarchies do construct discrete representations of higher-order categories, we hold that it is likely only a mechanism in service of *prediction* of feature changes across a *continuum of scales*. In this view, each interval along the continuum of perceptual processing in the ventral temporal stream and medial temporal lobe is simply specialized for learning the patterns of feature-change that occur over a particular interval of temporal and, as we will argue below, spatial scales [14–17].

Take, for example, an office. An office is not a discrete whole, but rather a collection of objects (e.g., walls, floor, computer, coffee mug, desk, chairs, bookshelves, among others). Neurally, an office as a whole tends to activate scene-sensitive regions such as PPA, while its constituents activate object-sensitive regions such as LOC. We assert that this difference in selectivity stems from the duration of time over which each region integrates the continuum of scales to register behaviorally-relevant changes in sensory features.

The behavioral significance of timescale can be easily illustrated in the office example depicted in Figure 1. The timescale over which our senses stably interact with individual objects, such as a coffee mug, is short; our eyes, as they make numerous saccades around the area, only process the mug for a matter of partial or full seconds, and our hands typically manipulate the mug for similar periods. On these short timescales, in which we must predictively guide our interactions with individual objects, it is necessary to maximally distinguish between, say, the cup and a stapler, hence the regional specialization of LOC for representing and detecting changes in patterns on this scale. By contrast, we may remain in the general office environment from several full minutes up to several hours, and even as the constituent objects repeatedly pass in and out of perception on their smaller timescales, we must integrate over all those features to extract a stable, longer-lasting context from the scene as a whole. In this latter case we would want to maximally distinguish this office scene context from others, such as a supermarket, a street corner, or a beach [18*]. This idea of a continuum of representational scales is shared by other recent work on Temporal Receptive Windows (TRWs) for representing narratives, whereby regions that integrate over large TRWs have highly divergent representations when individuals interpret narratives that differ in only a few words on a finer scale, but where those small

Figure 1



The brain represents change across a range of temporal scales. As our senses scan the environment, such as the office and the objects in the figure above (left panel), they integrate conceptual wholes over time from low-level features with activation level and specificity decaying logarithmically into the past (middle panel). Regions in the ventral temporal stream (VTS, right panel) are specialized for integrating over different spatiotemporal scales in order to maximally differentiate entities of a particular scale. In the example, LOC differentiates objects like the mug and the keyboard on a smaller timescale (shades of red), while PPA gradually integrates the same features into a larger-scale representation of the whole office (shades of blue) in order to differentiate it from other environmental contexts. These loci may not exist as anatomically discrete regions, but rather, as points on a continuum of scale-sensitivity seen in both the VTS and MTL (arrows in right panel). This scale continuum enables the efficient detection and prediction of patterns across all scales.

changes give rise to a very different meaning at larger scales [19].

Similarly, spatial relationships emerge from the stream of experience through time, extracted as structural regularities in features over different time scales. Referring back to the figure, the temporal stream of features representing objects in an office also includes retinotopic and, at least implicitly, spatiotopic information about the relative locations of those objects [20]. Thus, integrating over one time scale allows for extracting the typical configurations of objects in office scenes. Integrating over more time, we can represent the spatial relationships between larger regions of the environment.

Although the spectrum of scalar sensitivity in ventral temporal stream may itself be continuous, specific entities in our world can plausibly be expected to manifest more often within particular scale intervals [21], giving rise to anatomical clusters that tend to code for specific scales. Thus, the apparent representational discreteness in regions like ventral temporal cortex may simply result from our experience of scalar non-uniformities in our external environment, which give rise to groups of features that are structurally stable within the spatial and temporal ranges to which those localized regions are sensitive [22]. We alluded to this above, with the idea that the time scales over which we interact with objects can reliably just span seconds to minutes, whereas we are most often in locations for minutes to hours. Other components of our neural apparatus may link together features of experience that, because they have similar localization and extent, are likely to have arisen from the

same object or behaviorally-significant subdivision of the environment. This, too, is all in service of prediction because these shared features can be expected to co-occur again in future experiences.

It is important to stress that just because a region may be responding most strongly to features that would change over longer timescales (e.g., the PPA), this does not mean that it requires long integration times to represent that information during the normal perceptual process [23]. Instead, once a stable set of possible configurations for a scene has emerged over multiple experiences with scenes of that type, the regions representing the longer timescales will receive inputs from earlier visual regions and quickly activate representations that code for information relevant to the particular temporal scale that region tracks in order to provide maximal differentiation across experiences [24,25]. This initial representation can then be refined with additional feature input as the individual scans the environment.

This proposal suggests that regions specializing in differentiating experience integrated over longer and longer timescales will emerge through development as we gain more experience with various contexts at these different temporal scales. Some evidence in support of this claim exists such that children aged 7–11 have a relatively smaller PPA than adults, yet the LOC is no different [26]. This further suggests that this same principle would apply to earlier visual areas if researchers tracked neural representations earlier in development (e.g., LOC might exhibit relative differences between adults and younger children, perhaps ages 3–4).

A continuum of logarithmically-spaced scales enables efficient prediction

A proposal such as ours must not only describe a plausible predictive mechanism, but also explain how that particular mechanism would be favored by evolutionary constraints. As described above, life is driven to become maximally *efficient* given the computational costs and physical constraints of the processes that adaptively correlate it with the experiential environment. In other words, we do not have unlimited energy, neurons, and synapses to encode all of experience at full elaboration without becoming highly inefficient. Consequently, we must resolve this efficiency-utility trade-off by some means of abstraction [9,27]. It seems that one way the brain has addressed this problem is by extracting a scale-invariant representation from continuous experience. Also called Weber–Fechner law in perception, a scale-invariant representation implies a logarithmic encoding of the lag between feature change observations, with a corresponding fall-off in precision proportional to the size of the interval between those observations [28**].

Figure 1 illustrates the logarithmic scaling present as an individual scans a scene. Note how the representations of when the participant saw each object become less accurate the farther it happened in the past. Even though the figure focuses on tracking changes in *when* features activated, the same applies for *where*, as in the configural relationship between objects is more precise the closer they are.

Although the potential neural mechanisms that underlie this scale-free representation are beyond the scope of the present proposal [29], there exists a large and growing literature providing evidence for its existence, regardless of how it is implemented in the brain [16]. We instead review the reasons why this scale-invariant representation is an optimal solution for the efficiency-utility trade-off. First, it *allocates precision proportional to proximity and certainty*. The larger an interval, the more associative and causal uncertainty there will be as the uncertainty compounds with both temporal and spatial distance. Thus, a scheme that filters out information most likely to be uncertain/uninformative ‘noise’ in favor of more certain/informative ‘signal’ would be one that assigns precision in proportion to scale: a logarithmic representation. Second, this logarithmic representation *can accommodate associations over arbitrarily large scales*. Compared to many previous models that describe a finite memory buffer [30,31], scale-invariant models of memory and representation avoid the risk of neglecting associations with scales beyond the capacity of the buffer itself [16,32]. Third, a logarithmic representation of experience *exploits inter-scale dependencies to increase the cost-efficiency of building/learning the representations*. Here the algorithm harnesses stability from integrating over features of experience at larger scales to set expectations for features at

smaller scales, then the system can focus on only processing and learning maximally from deviations in this predicted stability. For example, as the ventral temporal cortex processes a scene in the woods, the recognition of the global scene properties as ‘woods-like’ by PPA would bias LOC toward interpreting small, green objects as leaves rather than, say, cans of Mountain Dew, and tall, thin objects as tree-trunks instead of streetlamps. This biasing sharply reduces the cost of processing the average situation, allowing the system to focus maximally on deviations from what was predicted.

From predictive coding in cortex to memory in the MTL

Our proposed scale-free spatiotemporal gradient of representation in the cortex dovetails nicely with theories of predictive coding and the free energy principle [33,34*]. In this paradigm, the brain constructs representations of complex and spatiotemporally extensive entities from simpler and more localized constituents, and in turn, these higher-level representations help to predict what will come next at the simpler, smaller level. The predictions are sent down the hierarchy, where they are squared against incoming information, generating error signals moving back up the hierarchy [33]. Thus, each link on the chain infers patterns of a given complexity/spatiotemporal scale, and each link also informs, and is informed by, its neighbors’ inferences. In this way, both finer-scale and coarser-scale context can be predicted from a stimulus at one certain scale — effectively allowing for inter-scale pattern-completion. However, this is still a temporally (and therefore spatially) local associative architecture. Each link in the chain is informed only by what its neighbors are identifying right *here* and *now*. They cannot, on their own, exploit associative information that is spatiotemporally removed from the present. They can rapidly and efficiently derive a rich picture spanning many scales from a sparse sample, while minimizing the energy it takes to represent that information, but that picture still only represents what is currently happening, along with the logarithmic gradient of recent experience leading up to this point. This graded representation has limited ability to extrapolate in detail across many experiences.

There is growing evidence that without a hippocampus (or as the MTL degrades with increasing dementia), a person’s ability to simulate the future becomes more and more generic and general [35]. This, along with a wealth of other data [36], points to the critical role of the MTL in both forming long-range associations and generating predictions of the future because the MTL is necessary for our ability to relive and draw on earlier experiences to generate predictions, and even to generate complex predictive value-maps to guide future actions [37]. In other words, the MTL allows us to mentally time travel [38] and, hence, flexibly navigate our past to make detailed predictions of the future across spatiotemporal scales, whereas the cortex can integrate across scales based on

the present, but only by generalizing so much that specificity is lost.

In order to accomplish this feat, we argue that the same scale free representations that emerge during continuous experience in the cortex are mirrored along the longitudinal axis of the MTL [15,39^{*}], though with the added architecture for quickly forming associations between sets of features *reinstated* from past experience [40]. For example, after experiencing $A \rightarrow B$ at one point and then $B \rightarrow C$ at another, the hippocampus makes the transitive inference between $A \rightarrow C$ possible because experiencing B the second time will reinstate A and bind it to C [41]. Consequently, the MTL is able to form associations between events that share features, but occur at completely different times. In functional imaging studies, the hippocampus, and particularly subfield CA1, has in fact been identified as central to such transitive-inferential associations [42]. These enhanced transitive associations mean the organism can not only access the logarithmically-compressed history of very recent and local experience, but that it can harness the MTL to reinstantiate experience from any significant point in the whole of time and space through which it has passed previously and bind it to the present experience. The organism thus maintains a representation of the world that becomes efficiently abstracted with increasing distance from here and now, yet nonetheless supports high representational density for the subset of remote events that are highly relevant despite their spatiotemporal distance. Through this process the organism is able to create a rich cognitive map of the world through the lens of experience [43].

We, therefore, interpolate between actually-experienced representations to create a complex model of the world that serves as an extended and flexible map. This map allows not only for highly-selective reaction to the current situation, which non-MTL cortical hierarchies can produce already, but also proactive pathfinding of multiple future courses of navigation and action based on previous experiences in related contexts [[44^{*},45]. Long-range associations can be enhanced further through replay of sequences of feature-states from the past or even *imagining* entirely new sequences constructed from those states. All of the reinstated states of the chain then become available once again for binding not only to the current, exogenous sensory input, but to *each other*. Therefore, long after their originating events have passed, large numbers of new predictive associations can be built between feature-state representations stored in memory. This highly developed, associative web implemented in the brain by the MTL and, in particular, the hippocampus is what that we most colloquially refer to with the word 'memory'; it serves a prospective function of correlating an organism's adaptive responses with the spatiotemporally extended regularities in the environment, thereby

creating the cognitive map over all experience and allowing us to be highly efficient *prediction machines*.

Time and space in memory

Neural time and space enable long-range associations

In expounding the virtues of harnessing the MTL to form long-range associations, we have neglected an important detail: how the brain decides what long-range associations to form. The short-range associative capability of unaided perceptual cortex has an inherent advantage in filtering for highly relevant associations between stimuli, since sensory inputs that are immediate neighbors in time are naturally related to one another. Pushing that associative frontier outwards in time would impair this quality-control, swamping the organism with a flood of long-distance associations, a much smaller proportion of which will improve the predictive signal enough to justify the added noise. To be adaptive, then, any long-range associative mechanism would have to filter candidate binding targets by some measure of relevance to the current context.

As argued above, the purpose of the spatiotemporal continuum in the brain is to detect and predict feature change at different scales. Detecting neural feature changes *also* provides the measure with which to calculate neural distance. For example, the more two experiences diverge in feature-change distance, the greater their separation in distance-based neural time and space [46,47]. Once such a neural distance measure is established, the probability that any particular past experience will be reactivated by, and then bound to, the current context can be made inversely proportional to its neural distance to the current context. The result is that experiences that are more *predictively relevant* are more likely to be associatively bound.

It is further possible to focus attention on a particular spatiotemporal scale more relevant for a particular task, which would change the anatomical locus of the distance metric along the longitudinal axis of the MTL. Evidence from gene-expression, anatomical, electrophysiological and fMRI data support the notion that there are functional long-axis scale gradients in the hippocampus, with anterior regions representing more global scales (e.g., larger place fields) and posterior regions representing finer scales (e.g., smaller place fields) [15,39^{*},48^{*},49]. Temporal and spatial memory are differentially impaired by damage to different positions on the long axis of the hippocampus, which suggests that this structural difference is also one of the typical scale of relevant feature change [50^{*}]. Further support for these long-axis gradients of scale comes from recent fMRI work demonstrating that neural distances between remembered real-life experiences spanning days to weeks and 100 meters to kilometers tracked these long-range time and space distances in the anterior hippocampus [51^{**}]. Recent work

also provides evidence that temporal and spatial representations are functionally intertwined in regions of the hippocampus, and that between-stimulus intervals in both dimensions are represented in a logarithmic fashion, suggesting Weber-Fechner scales, as predicted by the spatiotemporal continuum of feature change described above [17,51^{**},52^{**}].

Thus, space and time, which manifest in the spatiotemporal scales tracking feature change along the longitudinal axis of the MTL, may actually provide the critical distance metric that makes relevant long-range associations possible. These same spatiotemporal distance metrics may also provide a means for targeting and reinstating past experiences when more-specific content-based cues are lacking, by reconstructing their positions within the real-world dimension in which those experiences happened [53]. As stated above, long-range associations also allow for the construction of rich cognitive maps of experience, allowing for flexible prediction across spatiotemporal scales. In the following section we explore how even though space and time arise from the same distance metrics based on feature change, their informational content differs, which leads to distinct scenarios where they can be employed as useful memory cues.

Despite a common basis, neural time and space are informationally distinct

We have argued that neural time and space are both emergent properties of feature change along a continuum of scales, yet there is reason to believe they differ in information content, giving rise to priority differences in the memory system. To understand why this may be the case, we should first consider the characteristics of the real-world dimensions whose regularities these neural representations must emulate in order to serve their adaptive function. The most obvious concerns the number of subordinate dimensions; external time has one, while external space has three, affording a geometrically larger set of associative possibilities between events. Experiences in space can be interrelated in complex topological networks, while experiences in time can only occupy earlier or later positions on a line. Equally important, the flow of time is also unidirectional, whereas it is possible to revisit spatial locations.

These qualitative differences have major consequences for us as we interact with the world. Although we can return to the same point in real space during many different experiences, each time arriving and departing along a different trajectory in both real and perceptual space, we have no corresponding ability to revisit a point in real time. Lacking in these repetitions, time appears inherently less stable than space from the perspective of a living system. However, the representational architecture of neural space appears to leverage the additional

information, incorporating a variety of intermediate constructs that have no apparent counterparts in the computation of neural time. These include optic flow [16], ocular motion [54], direction [55,56], vestibular cues [57^{*}], as well as a number of more integrative mechanisms including boundary cells [58], speed cells [59], place cells [60], and grid-like neuronal firing fields that map to evenly distributed spatial geometries and also follow the same coarse to fine representational continuum along the longitudinal axis of the MTL [61,62,63^{*}]. Each of these captures some contingency between change in one subordinate dimension and change in another, and functions to modulate and reinforce the underlying distance-based representation depending on those constraints. By incorporating this additional between-dimension contingency information over and above the more basic contextual-distance measures, the neural representations of space thus provide stronger and more reliable retrieval cues than those of time.

It is informative to consider the case of a sharp boundary, such as the transition a person experiences moving indoors to outdoors. Such transitions are marked by large changes in feature-space, which give rise to large changes in the spatiotemporal context used to guide predictions. These large feature changes also likely provide the anchor points to help identify event boundaries at particular scales [4]. This scale specificity for tracking event boundaries is underlined by the fact that just because you transitioned from inside to outside your office building it does not mean you have left the context of your university, which can still provide useful predictions, such as seeing students and colleagues, at larger and longer spatiotemporal scales. Finally, just because modern life may tend to have more sharp transitions than the past (i.e., with fewer buildings, elevators, or motorized vehicles, especially air travel), this does not mean that the general neural architecture we propose here has actually changed over such a short evolutionary timescale. Instead, it may be that these sharp changes in spatiotemporal features are one way in which modern life induces stress in neural processing.

Conclusions

Here we have recast the neural machinery of perception, memory, and prediction, away from the classic narrative of a set of discrete brain regions whose main role is to represent higher-order categories, toward a continuum of scales whose aggregate function is to predict future changes in streams of elemental sensory features based on the history of those streams. The foundation of this system is a perceptual continuum in cortex, where each point on the continuum integrates over, and maximally responds to, a particular spatiotemporal scale of feature-change. Each integrator cross-talks with its higher-scale and lower-scale neighbors to leverage predictive coding for greater classification accuracy and efficiency.

Together, the full continuum provides a comprehensive representation of the trans-scale immediate context and a prediction of its very immediate future.

Nonetheless, without the contribution of the medial temporal lobe, particularly the hippocampus, this prediction is limited to a tiny slice of the present stream of experience, decaying logarithmically in both forward and backward directions. Neural time and space provide a distance metric derived from feature change across these spatiotemporal scales that the MTL can harness to identify *relevant* similarity of past experiences to the present. This enables reinstatement of remote contexts, which can be bound to the present experience to support long-range associations or to enrich the richness of predictions from the present. The long-range associations ultimately lead to the construction of cognitive maps across the entire spectrum of feature change, integrating over all life's experiences.

In this sense, semantic knowledge would constitute experiences embedded in a map that has been so thoroughly elaborated that they 'transcend' neural time and space entirely to become universally-applicable concepts and facts, since there is no longer anything to anchor them to any specific temporal or spatial contingency. This would render the classically sharp division between episodic and semantic memory a difference of extremity, not of kind. Instead, episodic and semantic memory both emerge due to the same extraction of regularities in the continuum of spatiotemporal feature change across experiences [64].

Implications for future research

This synthesis has a number of experimental ramifications. First, we should watch for opportunities to resolve apparent contradictions in functional results that stem from expectations that perceptual brain regions primarily respond to specific, discrete categories. The example of [65] is instructive. If the PPA represents only place identity, then the finding that it also responds to material category requires radical reevaluation of the first hypothesis; but if PPA instead recognizes feature-change at a scale common to both place and material, there is no contradiction.

Second, by focusing inquiry on the function of isolated loci in perceptual cortex and medial temporal lobe, we may be missing more important trends happening across continua. In order to mitigate this, future work in this area should record and analyze data from across a much wider range of scales in both space and time [51**]. This is facilitated by the proliferation of life-logging technology which permits persistent collection of multi-modal data (digital imaging, accelerometry, GPS, among others) across the naturalistic settings of the subjects' own lives [66].

Third, in the realm of computational modeling, particular attention should be given to models with logarithmic representations of spatiotemporal scale and distance following the Weber-Fechner law [67,68*].

Fourth, if the apparent discreteness of perceptual loci is not endogenous, but results from actual scalar non-uniformities in the structure of the external sensory environment, this raises the interesting question of whether organisms learning in environments with artificially-distorted scalar distributions would come to recognize alternate dimensions of feature change that do not readily correspond with the familiar labels we currently utilize to describe our world, similar to the way brain regions for vision are co-opted for other forms of representation in congenitally blind individuals [69].

Fifth, if higher-order perceptual integrators depend on smaller-scale integrators acting in concert with the hippocampus for learning at their own scale, this suggests that earlier visual areas should follow the same sequential pattern of development as LOC and PPA. Comparing relative differences in neural representations in these areas between adults and younger children would help to determine if this is the case.

Finally, to our knowledge, no study to date has been performed with the explicit goal of comparing the representational boundaries of time and space. An experimental paradigm that intentionally manipulates temporal contexts in order to enhance their space-like characteristics (such as the recurrence of temporal 'landmarks') and then compares the resulting patterns of neural activations with those based on true spatial contexts would help to answer this.

Conflict of interest statement

Nothing declared.

Acknowledgements

We would like to thank Adam W. Fenton, Ryan P. Kirkpatrick, and Emily R. Weichart for comments and suggestions on the manuscript. This work was supported by the National Science Foundation (1631403) and the Rudi Schulte Research Institute. Despite financial support, neither funding agency had a specific role in writing or the decision to submit the article.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Tulving E: **Episodic and semantic memory**. In *Organization of Memory*. Edited by Tulving E, Donaldson W. Academic Press; 1972:381-402.
 2. Humphreys MS, Wiles J, Dennis S: **Toward a theory of human memory: data structures and access processes**. *Behav Brain Sci* 1994, **17**:655-667.
 3. Neath I, Surprenant A: *Human Memory: An Introduction to Research, Data, and Theory*. edn 2. Wadsworth; 2003.

4. Zacks JM, Swallow KM: **Event segmentation**. *Curr Direct Psychol Sci* 2007, **16**:80-84.
5. DiCarlo JJ, Zoccolan D, Rust NC: **How does the brain solve visual object recognition?** *Neuron* 2012, **73**:415-434.
6. Grill-Spector K, Kourtzi Z, Kanwisher N: **The lateral occipital complex and its role in object recognition**. *Vision Res* 2001, **41**:1409-1422.
7. Epstein R, Kanwisher N: **A cortical representation of the local visual environment**. *Nature* 1998, **392**:598-601.
8. Grill-Spector K, Weiner KS: **The functional architecture of the ventral temporal cortex and its role in categorization**. *Nat Rev Neurosci* 2014, **15**:536-548.
9. Still S, Sivak DA, Bell AJ, Crooks GE: **Thermodynamics of prediction**. *Phys Rev Lett* 2012, **109**:120604.
10. Wolpert D, Kolchinsky A: *Observers as Systems That Acquire Information to Stay Out of Equilibrium*. In Banff; 2016.
11. Rovelli C: **Meaning=Information+Evolution**. arXiv:161102420 [physics] [Internet]. 2016 Nov [cited 2017 Jan 30]; Available from: <http://arxiv.org/abs/1611.02420>.
12. Gershman SJ, Moore CD, Todd MT, Norman KA, Sederberg PB: **The successor representation and temporal context**. *Neural Comput* 2012, **24**:1553-1568.
13. Rankin CH, Abrams T, Barry RJ, Bhatnagar S, Clayton D, Colombo J *et al.*: **Habituation revisited: an updated and revised description of the behavioral characteristics of habituation**. *Neurobiol Learn Memory* 2009, **92**:135-138.
14. Amoruso L, Couto B, Ibáñez A: **Beyond extrastriate body area (EBA) and fusiform body area (FBA): context integration in the meaning of actions**. *Front Hum Neurosci* 2011, **5**:124.
15. Poppenk J, Evensmoen HR, Moscovitch M, Nadel L: **Long-axis specialization of the human hippocampus**. *Trends Cogn Sci* 2013, **17**:230-240.
16. Howard MW, Eichenbaum H: **The hippocampus, time, and memory across scales**. *J Exp Psychol Gen* 2013, **142**:1211-1230.
17. Buzsáki G: **Our skewed sense of space**. *Science* 2015, **347**:612.
18. Livne T, Bar M: **Cortical integration of contextual information across objects**. *J Cogn Neurosci* 2016, **28**:948-958.
Discussion of the integrative dynamics that infer higher-order wholes from contextual feature-streams. Discusses spatial coherence as a 'second-order contextual property' in comparison to the more basic properties of contextual co-occurrence between features in these streams.
19. Yeshurun Y, Nguyen M, Hasson U: **The butterfly effect: amplification of local changes along the temporal processing hierarchy**. *bioRxiv* 2017 Jan.
20. Golomb JD, Kanwisher N: **Higher level visual cortex represents retinotopic, not spatiotopic, object location**. *Cereb Cortex* 2012, **22**:2794-2810.
21. Sreekumar V, Dennis S, Doxas I, Zhuang Y, Belkin M: **The geometry and dynamics of lifelogs: discovering the organizational principles of human experience**. Edited by Balasubramaniam R. 2014, **9**:e97166.
22. Ferber S, Humphrey GK, Vilis T: **The lateral occipital complex subserves the perceptual persistence of motion-defined groupings**. *Cereb Cortex* 2003, **13**:716-721.
23. O'Connell TP, Walther DB: **Dissociation of salience-driven and content-driven spatial attention to scene category with predictive decoding of gaze patterns**. *J Vision* 2015, **15**:20-0.
24. Güçlü U, van Gerven MAJ: **Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream**. *J Neurosci* 2015, **35**:10005.
Supports a complexity gradient in the ventral temporal stream, demonstrated by layering in deep neural network models of that region.
25. Cichy RM, Pantazis D, Oliva A: **Similarity-based fusion of MEG and fMRI reveals spatio-temporal dynamics in human cortex during visual object recognition**. *Cereb Cortex* 2016, **26**:3563-3579.
26. Golarai G, Ghahremani DG, Whitfield-Gabrieli S, Reiss A, Eberhardt JL, Gabrieli JDE *et al.*: **Differential development of high-level visual cortex correlates with category-specific recognition memory**. *Nat Neurosci* 2007, **10**:512-522.
27. Levy WB, Ju H, Baxter RA, Colbert CM: **Controlling information flow and energy use via adaptive synaptogenesis**. *2016 Annual Conference on Information Science and Systems (CISS)*. 2016:535-538.
28. Howard MW, Shankar KH: **Neural scaling laws for an uncertain world**. arXiv:160704886 [physics, q-bio] [Internet], submitted [cited 2017 Jan 30]; Available from: <http://arxiv.org/abs/1607.04886>.
Describes many cases where the scale-invariant Weber-Fechner law is found to operate in perceptual mechanisms, and explains how this permits the brain to represent information from an uncertain world without prior assumptions.
29. Howard MW, MacDonald CJ, Tiganj Z, Shankar KH, Du Q, Hasselmo ME *et al.*: **A unified mathematical framework for coding time, space, and sequences in the hippocampal region**. *J Neurosci* 2014, **34**:4692-4707.
30. Raaijmakers JG, Shiffrin RM: **Search of associative memory**. *Psychol Rev* 1981, **88**:93-134.
31. Davelaar EJ, Goshen-Gottstein Y, Ashkenazi A, Haarmann HJ, Usher M: **The demise of short-term memory revisited: empirical and computational investigations of recency effects**. *Psychol Rev* 2005, **112**:3-42.
32. Sederberg PB, Howard MW, Kahana MJ: **A context-based theory of recency and contiguity in free recall**. *Psychol Rev* 2008, **115**:893-912.
33. Friston K: **The free-energy principle: a rough guide to the brain?** *Trends Cogn Sci* 2009, **13**:293-301.
34. Gordon N, Koenig-Robert R, Tsuchiya N, van Boxtel JJ, Hohwy J: **Neural markers of predictive coding under perceptual uncertainty revealed with Hierarchical Frequency Tagging**. Edited by Stephan KE. 2017, **6**:e22749.
Utilizes hierarchical frequency tagging of stimuli in order to detect signatures of predictive-coding at work in the ventral temporal stream, which we argue is facilitated by gradients of spatiotemporal scale.
35. Addis DR, Schacter D: **The hippocampus and imagining the future: where do we stand?** *Front Hum Neurosci* 2012, **5**:173.
36. Buckner RL: **The role of the hippocampus in prediction and imagination**. *Annu Rev Psychol* 2010, **61**:27-48.
37. Stachenfeld KL, Botvinick MM, Gershman SJ: **The hippocampus as a predictive map**. *bioRxiv* 2016 Dec.
38. Polyn SM, Sederberg PB: **Brain rhythms in mental time travel**. *NeuroImage* 2014, **85**(Part 2):678-684.
39. Strange BA, Witter MP, Lein ES, Moser EI: **Functional organization of the hippocampal longitudinal axis**. *Nat Rev Neurosci* 2014, **15**:655-669.
In-depth review of hippocampal longitudinal specialization. Offers a model of a functional gradient of scale superimposed on locally discrete regional divisions defined by gene expression and connectivity. It is consistent with models of a continuum of scale-sensitivity, while accommodating the possibility of discrete underlying mechanisms.
40. Howard MW, Kahana MJ: **A distributed representation of temporal context**. *J Math Psychol* 2002, **46**.
41. Howard MW, Jing B, Rao VA, Provyn JP, Datey AV: **Bridging the gap: transitive associations between items presented in similar temporal contexts**. *J Exp Psychol Learn Memory Cogn* 2009, **35**:391-407.
42. Schlichting ML, Zeithamova D, Preston AR: **CA1 subfield contributions to memory integration and inference**. *Hippocampus* 2014, **24**:1248-1260.
43. McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser M-B: **Path integration and the neural basis of the 'cognitive map'**. *Nat Rev Neurosci* 2006, **7**:663-678.
44. Hasselmo ME: **A model of episodic memory: mental time travel along encoded trajectories using grid cells**. *Neurobiol Learn Memory* 2009, **92**:559-573.

Proposes a model that integrates place, grid, direction, boundary, speed and arc-length cells to build a complex representation of space and the actions and events that are associated with points in it. This representation can support the simulation of spatiotemporal trajectories, both in episodic recall and prospecting.

45. Shankar KH, Singh I, Howard MW: **Neural mechanism to simulate a scale-invariant future.** *Neural Comput* 2016, **28**:2594-2627.
46. Manns JR, Howard MW, Eichenbaum H: **Gradual changes in hippocampal activity support remembering the order of events.** *Neuron* 2007, **56**:530-540.
47. Ezzyat Y, Davachi L: **Similarity breeds proximity: pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity.** *Neuron* 2014, **81**:1179-1189.
48. Javadi A-H, Emo B, Howard LR, Zisch FE, Yu Y, Knight R *et al.*: **Hippocampal and prefrontal processing of network topology to simulate the future.** *Nat Commun* 2017, **8**:14652.
Characterizes spatial representations in medial temporal lobe as fundamentally topological rather than geometric, according well with the characterization of neural space that is constructed from sensory content instead of preceding it as a scaffold for that content.
49. Brunec IK, Bellana B, Ozubko JD, Man V, Robin J, Liu Z-X *et al.*: **Differential spatiotemporal representations along the hippocampal long axis in humans.** *bioRxiv* 2017 Aug.
50. Barker GRI, Banks PJ, Scott H, Ralph GS, Mitrophanou KA, Wong L-F *et al.*: **Separate elements of episodic memory subserved by distinct hippocampal-prefrontal connections.** *Nat Neurosci* 2017 Feb, **20**:242-250.
Demonstrated that inactivating connections between mPFC and either dorsal or intermediate hippocampus (CA1) in rats impaired memory differently: dorsal inactivation selectively impaired temporal order judgments while intermediate inactivation impaired more spatial memory. This supports the idea that representations of neural time and space may differ in scale along the longitudinal axis of the hippocampus.
51. Nielson DM, Smith TA, Sreekumar V, Dennis S, Sederberg PB: **Human hippocampus represents space and time during retrieval of real-world memories.** *Proc Natl Acad Sci U S A* 2015, **112**:11078-11083.
Provides evidence that the anterior hippocampus represents both space and time at very large scales. The results support the idea that neural space and time are intertwined in the hippocampus and that there is a gradient of small-scale to large-scale representations along the longitudinal axis of the hippocampus.
52. Deuker L, Bellmund JL, Navarro Schröder T, Doeller CF: **An event map of memory space in the hippocampus.** Edited by Davachi L. 2016, **5**:e16534.
Provides evidence that temporal and spatial aspects of episodic memory are supported by the same coding mechanism in anterior hippocampus, with implications for the construction of cognitive maps.
53. Curran T, Friedman WJ: **Differentiating location- and distance-based processes in memory for time: an ERP study.** *Psychonomic Bull Rev* 2003, **10**:711-717.
54. Golomb JD, Albrecht AR, Park S, Chun MM: **Eye movements help link different views in scene-selective cortex.** *Cereb Cortex* 2011, **21**:2094-2102.
55. Taube JS: **Head direction cells and the neurophysiological basis for a sense of direction.** *Progr Neurobiol* 1998, **55**:225-256.
56. Taube JS: **New building blocks for navigation.** *Nat Neurosci* 2017, **20**:131-133.
57. Aghajani ZM, Acharya L, Moore JJ, Cushman JD, Vuong C, Mehta MR: **Impaired spatial selectivity and intact phase precession in two-dimensional virtual reality.** *Nat Neurosci* 2015, **18**:121-128.
A virtual reality study in rats consistent with the notion that rich information from multiple sensory sources is necessary to support spatial coding. Barring that informational richness, only temporal coding is found, as expected by our observations of the differences between the physical dimensions of time and space.
58. Lever C, Burton S, Jeewajee A, O'Keefe J, Burgess N: **Boundary vector cells in the subiculum of the hippocampal formation.** *J Neurosci* 2009, **29**:9771-9777.
59. Kropff E, Carmichael JE, Moser M-B, Moser EI: **Speed cells in the medial entorhinal cortex.** *Nature* 2015, **523**:419-424.
60. O'Keefe J, Dostrovsky J: **The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat.** *Brain Res* 1971, **34**:171-175.
61. Doeller CF, Barry C, Burgess N: **Evidence for grid cells in a human memory network.** *Nature* 2010, **463**:657-661.
62. Agarwal G, Stevenson IH, Berenyi A, Mizuseki K, Buzsáki G, Sommer FT: **Spatially distributed local fields in the hippocampus encode rat position.** *Science* 2014, **344**:626-630.
63. Kyle C, Smuda D, Hassan A, Ekstrom A: **Roles of human hippocampal subfields in retrieval of spatial and temporal context.** *Behav Brain Res* 2015, **278**:549-558.
Finds that, while there are hippocampal differences in multivariate patterns of temporal and spatial contexts, no hippocampal subfields appear to be specialized for either, supporting the notion that the differences between neural time and space arise because of structural differences in the exogenous dimensions on which these representations are based, not because of divergent coding mechanisms.
64. Howard MW, Shankar KH, Jagadisan UKK: **Constructing semantic representations from a gradually-changing representation of temporal context.** *Topics Cogn Sci* 2011, **3**:48-73.
65. Jacobs R, Baumgartner E, Gegenfurtner K: **The representation of material categories in the brain.** *Front Psychol* 2014, **5**:146.
66. Dennis S, Yim H, Sreekumar V, Evans N, Garret P, Sederberg P: **A hierarchical Bayesian model of memory for when based on experience sampling data.** *Cogn Sci* 2017. Submission # 1333.
67. Howard MW, Shankar KH, Aue WR, Criss AH: **A distributed representation of internal time.** *Psycho Rev* 2014, **122**:24-53.
68. Howard MW, Shankar KH, Aue WR, Criss AH: **A distributed representation of internal time.** *Psychol Rev* 2015, **122**:24-53.
Description of a coding scheme for scale-invariant representations of temporal history that can be applied to all sensory modalities and corresponds with a range of observed behavioral phenomena.
69. Lane C, Kanjlia S, Omaki A, Bedny M: **"Visual" cortex of congenitally blind adults responds to syntactic movement.** *J Neurosci* 2015, **35**:12859.