

REPORT DOCUMENTATION PAGE					<i>Form Approved OMB No. 0704-0188</i>	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Services and Communications Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.						
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.						
1. REPORT DATE (DD-MM-YYYY) 31-07-2012		2. REPORT TYPE FINAL			3. DATES COVERED (From - To) 30-4-2010 31-7-20	
4. TITLE AND SUBTITLE A distributed representation of remembered time					5a. CONTRACT NUMBER	
					5b. GRANT NUMBER FA9550-10-1-0149	
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Howard, Marc W.					5d. PROJECT NUMBER	
					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Boston University, 2 Cummington St, Boston MA 02215					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office for Scientific Research/RSL 875 N. Randolph Street Arlington, VA 22203					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-OSR-VA-TR-2012-1121	
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution A: Approved for public release.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT The overarching objective of the proposed research was to develop an extension of the temporal context model to enable the description of a broad variety of phenomena across various subfields of cognitive psychology by integrating a representation of the time at which stimuli were experienced. This was accomplished. We published a paper in Neural Computation that describes such a model and applied it to problems in episodic memory, timing and classical conditioning. A more detailed version of this model, described in a paper revised for Psychological Review, extends it to a wider range of phenomena by introducing a translation operator allowing for the construction of trajectories of predicted future states, and jumping-back-in-time to allow for an account of contiguity effects in episodic memory. In that paper we applied the model to problems from episodic memory, working memory, as well as second-order conditioning problems from trace conditioning. We met the near-term goals of generating simulation code for the model of timing, with three different approaches to simulating the equations implemented in the R program language.						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Howard, Marc W.	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) 617 353 8448	

Reset

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Final Report for FA9550-10-1-0149, A distributed representation of space and time

Marc W. Howard

Department of Psychology and Center for Memory and Brain
Boston University

Prepared 31-07-2012

Abstract

The overarching objective of the proposed research was to develop an extension of the temporal context model to enable the description of a broad variety of phenomena across various subfields of cognitive psychology by integrating a representation of the time at which stimuli were experienced. This was accomplished. We published a paper in *Neural Computation* that describes such a model and applied it to problems in episodic memory, timing and classical conditioning. A more detailed version of this model, described in a paper revised for *Psychological Review*, extends it to a wider range of phenomena by introducing a translation operator allowing for the construction of trajectories of predicted future states, and jumping-back-in-time to allow for an account of contiguity effects in episodic memory. In that paper we applied the model to problems from episodic memory, working memory, as well as second-order conditioning problems from trace conditioning. We met the near-term goals of generating simulation code for the model of timing, with three different approaches to simulating the equations implemented in the R program language. One is based on simulating the time series of inputs as a series of delta functions. The second approximates the representation of history by implementing a partial differential equation it obeys. The third, which is most general but subject to approximation errors explicitly calculates the operator that inverts the Laplace transform. The sensitivity to noise of the model was developed in the *Neural Computation* paper. We also made progress towards the longer-term goals of the proposal. Although not yet published, we have analyzed the properties of the model as applied to semantic memory. Although we have not conducted the simulations, we successfully secured funding from the NSF to pursue that longer-term goal.

FA9550-10-1-0149 supported the development of a quantitative model of stimulus history and applied it to numerous phenomena in cognitive psychology. That model developed a mathematical formalism in which a stimulus history $f(\tau)$ is encoded by means of the Laplace transform, and an approximate inversion.

The research proposed several sub-projects grouped roughly into near-term and medium-term projects. Here we summarize the final status of these sub-projects. After this brief summary, we

describe the scientific content of the output of the research project in considerably more detail.

Near-term projects

Near-term projects were anticipated to be completed in the first year of funding. Both of these were completed as anticipated.

1. *Development of a simulation library for remembered time.* The model started with a mathematical formulation. In order to facilitate comparison of the equations to behavioral and neural results, we proposed to develop a software simulation library to implement the equations. As it turns out we developed three versions of the library, each corresponding to a different way of implementing the equations, in the R programming language. We have provided these libraries to other investigators on request to facilitate widespread use of the model in scientific investigation. We have also facilitated the development of a Python library based on our R code by the Dennis lab at Ohio State.

2. *Analytic study of recency and contiguity effects.* The recency effect and the contiguity effect refer to two fundamental aspects of episodic memory. The recency effect refers to the finding that, all other things being equal, stimuli experienced more recently in the past are better remembered than stimuli experienced less recently in the past (Murdock, 1962; Glenberg et al., 1980). The contiguity effect refers to the finding that, all other things being equal, when one stimulus comes to mind, the next stimulus that comes to mind is likely to have been experienced close in time to the first stimulus (Kahana, 1996). Both of these phenomena have been observed over a variety of time scales (Glenberg et al., 1980; Howard & Kahana, 1999; Howard, Youker, & Venkatadass, 2008; Unsworth, 2008) suggesting a common mechanism that is not based on traditional ideas of short-term memory.

We proposed to analyze these findings in the context of our model of remembered time. The basic idea is that the current history serves as a cue to initiate recall, leading to the recency effect. Contiguity was hypothesized to happen when a stimulus is remembered causes recovery of the state of history when it was initially experienced. Because the representation of history is scale-free, it should account for both findings over a variety of time scales.

The model's treatment of the recency effect was first addressed in Shankar and Howard (2012). Recency worked as expected. The account of contiguity effect was slightly more involved and was presented in Howard, Shankar, Aue, and Criss (In preparation). Interestingly, the model ends up predicting violations of strict scale-invariance in the contiguity effect that appear consistent with experimentally-observed findings. Because of the centrality of contiguity in theories of episodic memory, we were pleased to contribute to two unanticipated side projects that extended this goal beyond our initial expectations. One was completion of revisions of a paper confirming neurophysiological predictions of the account of contiguity based on recovery of temporal information (Howard, Viskontas, Shankar, & Fried, In press). In that paper, the pattern of activity in ensembles of neurons in the human MTL showed evidence for a "jump back in time" as we had anticipated. In a second paper, we helped design and analyze an empirical study that developed a new experimental methodology for measuring the contiguity effect behaviorally in human subjects (Kılıç, Criss, & Howard, In press). Because that study established a causal relationship between the presentation of the cue and the recovered memories, it cannot be accounted for by several competing explanations of the contiguity effect (Davelaar, Goshen-Gottstein, Ashkenazi, Haarmann, & Usher, 2005; Grossberg & Pearson, 2008; Farrell, 2012).

3. *Sensitivity to noise.* This subproject was designed to evaluate the model’s resistance to stochastic noise in its inputs. This is an important topic to investigate because if the model is too sensitive to disruptions, it would be unable to represent stimulus history over the long periods of time necessary to account for the behavioral phenomena under investigation. These analyses were conducted and reported in detail in Shankar and Howard (2012). Briefly, the representation of stimulus history is surprisingly robust to stochastic disruptions in the input patterns. We further examined the detailed response of the model to signals with various spectral signatures in Shankar and Howard (submitted).

Medium-term projects

Medium-term projects were ones that would not be likely to be completed within a single year. We were hopeful to complete these within the two year window. As it turned out, we successfully completed one of these subprojects and made strong progress on the second.

Temporal mapping in trace conditioning. Previous work in the mathematical modeling of memory focused either on human memory experiments in which subjects learn lists of words, or on conditioning experiments in which rats learn to respond appropriately to stimuli presented on various schedules. This parochial approach seems to us to be poorly suited to the goal of developing a general theory of memory that applies across species. With this in mind, we proposed to describe phenomena from the temporal mapping literature (e.g., Cole, Barnet, & Miller, 1995), a conditioning paradigm in which rats learn to integrate temporal relationships between stimuli (see below for more detail).

We describe our treatment of temporal mapping in some detail below and in considerably more detail in Howard et al. (In preparation). Briefly, the solution turned out to require major advances in the mathematical framework. First, as anticipated, we required the same “jump back in time” that we used to account for the contiguity effect in episodic memory studies, suggesting a deep analogy between human list learning and animal classical conditioning studies. Second, we needed to construct a “translation operator” that enables the model to “play forward” anticipated future states of the world using the current state of history. This translation operator (described below) also, fortuitously enabled an account of other behavioral phenomena, namely interval timing across scales and the “time-left procedure” from animal conditioning studies (see Howard et al., In preparation, for details).

Modeling the development of semantic representations in higher-order statistical environments. Episodic memory refers to the ability to remember specific instances from one’s life. Episodic memories are rich in detail, as if the rememberer is reexperiencing the event. In contrast semantic memory reflects general verbalizable knowledge that does not depend on reexperiencing the specific learning event. For instance, the question “what did you have for breakfast” usually evokes an episodic memory in which one in some sense relives their most recent breakfast as part of answering the question. In contrast, “what is the capital of Vermont” does not usually evoke a vivid recollection of the moment when that fact was learned. We extensively applied the mathematical framework to phenomena from episodic memory experiments. The goal of this subproject was to extend the formalism to account for semantic memory phenomena as well. Although not yet published, we have solved several problems that will be necessary for describing semantic memory.

The basic idea is that meaning in semantic memory is developed largely by learning the temporal roles played by different stimuli. For instance, consider the novel word FLOOB. Prior to

experience, there is no information about the meaning of FLOOB. However, after a single exposure in the sentence “The baker reached into the oven and pulled out the FLOOB,” a tremendous amount of information about the meaning of FLOOB is available to the learner. The basic hypothesis is that we can use the current state of temporal history to generate a prediction. This prediction can be used to inform the meaning of stimuli that are experienced (Shankar, Jagadisan, & Howard, 2009; Howard, Shankar, & Jagadisan, 2011).

Previous work used this general approach to learn from a much simpler representation of temporal context. Our long-term goal is to utilize the richer representation of temporal history we developed to fulfill the same role. Previously we established that the version of the model based on temporal context can learn perfectly any language that can be described by bigram transition probabilities (Shankar et al., 2009). That is all models in which the identity of the next symbol can be deduced from only the prior symbol. Unpublished analyses have shown that the expanded model should be able to account for lagged bigram languages. That is, languages in which the next symbol can be deduced from independently combining the identity of the previous symbol and the identity of the symbol before that and so on. This limitation is actually a huge advantage over trying to learn N-gram languages which would require an astronomical number of observations to learn the probability of non-independent combinations of symbols. Moreover, the lags become less precise the further in the past one is considering. That is, the model distinguishes the symbol at lag 1 and the symbol at lag 2 with more temporal resolution than it does the symbol at lag 11 and the symbol at lag 12. Subject to this constraint, the model can capture long-range correlations among symbols. We suspect that, because language includes constraints over a wide variety of time scales, this property constitutes an important advance over existing computational models of semantic memory (e.g., Landauer & Dumais, 1997; Griffiths, Steyvers, & Tenenbaum, 2007).

As a practical implementation matter, we have determined that the generalization from temporal context to temporal history is not a major obstacle because of the fact that history can be composed from states of temporal context and the linearity of the operator making that construction. The major open problem to be solved in describing language is polysemy. In practice, language cannot simply be described as a lagged bigram language. The problem is that polysemous symbols—words with multiple meanings—play different roles in different contexts. This prevents the statistics from being expressed as a lagged bigram language: if what follows from the word FLY depends on whether it’s in a discussion of baseball or a discussion of insects, this prevents us from treating the effect of that symbol independently of the other symbols nearby. The solution is to generate a new set of symbols in which polysemous meanings are distinguished by their context. This amounts to a blind source separation problem, which is relatively well-understood. We have secured funding from NSF (BCS-1058937) to pursue this line of research and construct a large-scale computational model from naturally-occurring text.

Detailed scientific description of results

Taken as a whole, the set of behavioral findings accommodated in this mathematical framework reflect a major advance towards a satisfactory theory of memory. Here we describe the model and the results in more detail so that the interested reader can appreciate these deep connections between fields. By means of introduction, we first briefly recap the formalism of the model and the behavioral simulations. These are described in more detail elsewhere (Shankar & Howard, 2012; Howard et al., In preparation).




Figure 1. Schematic illustrating the construction of $\mathbf{T}$ from $\mathbf{t}$ and the role of $\mathbf{M}$. **a.** In this schematic, the stimulus dimensions (columns) have been ordered with the sequence in which the stimuli were presented. Dark values correspond to higher levels of activation. Different rows correspond to different values of s . Each row of $\mathbf{t}$ represents the stimuli as an exponential function. The operator $\mathbf{L}_k^{-1}$ takes $\mathbf{t}$ into $\mathbf{T}$ at each moment. Stimuli presented in the recent past have a concise representation; stimuli further in the past have overlapping representations. For visual clarity, values in $\mathbf{T}$ have been scaled to their maximum value. **b.** At each moment, the current value of $\mathbf{f}$ is used to update the current value of $\mathbf{t}$. At each moment, the current value of $\mathbf{t}$ is used to update $\mathbf{T}$, which is, in turn used to update $\mathbf{M}$. The current value of $\mathbf{M}$ and the current state $\mathbf{T}$ are used to generate a prediction $\mathbf{p}$.

Representing the past

Given a stimulus history $\mathbf{f}(\tau)$ leading up to the present moment, the goal is to represent that function using only operations local in time. This can be accomplished by using a set of leaky integrators $\mathbf{t}(s)$ that each obey the differential equation

$$\frac{d\mathbf{t}}{d\tau} = -s\mathbf{t} + \mathbf{f}(\tau), \quad (1)$$

where τ is physical time. The solution to Eq. 1 is just the Laplace transform of $\mathbf{f}(\tau' < \tau)$:

$$\mathbf{t}(\tau) = \int_{-\infty}^{\tau} e^{-s(\tau-\tau')} \mathbf{f}(\tau') d\tau' \quad (2)$$

This means that $\mathbf{t}(s)$ contains all of the information in the stimulus up to that point. Armed with this insight, we can construct an approximation of the inversion of the Laplace transform. It turns out that a method attributable to Post (1930) has just the properties we want. Referring to the approximation operator for integer k as $\mathbf{L}_k^{-1}$, we write

$$\mathbf{T}(\tau)^* \equiv \mathbf{L}_k^{-1} \mathbf{t}(s) \quad (3)$$

$$= \frac{-1^k}{k!} s^{k+1} \mathbf{t}^{(k)}(s), \quad (4)$$

where $\mathbf{t}^{(k)}(s)$ is the k th derivative of $\mathbf{t}$ with respect to s . The approximation approaches the inverse Laplace transform as $k \rightarrow \infty$. $\mathbf{T}(\tau)^*$ approximates the stimulus history as a function of internal time τ . Figure 1a provides a schematic summary of the construction of $\mathbf{T}(\tau)^*$ from $\mathbf{t}(s)$ at any given moment.

Predicting the future

At each moment in time, we cannot only remember the past, but use history to predict the stimulus that will be presented at the next moment as well. This can be accomplished by comparing

the current state of $\mathbf{T}$ to the previous state of $\mathbf{T}$ in which the stimulus was experienced. Following TCM (Howard & Kahana, 2002; Howard, Fotedar, Datey, & Hasselmo, 2005; Sederberg, Howard, & Kahana, 2008), we store a simple outer product between the currently-experienced stimulus and the current state of $\mathbf{T}$ at each moment and store it in a tensor $\mathbf{M}$:

$$\frac{d\mathbf{M}}{d\tau} = |\mathbf{f}(\tau)\rangle\langle\mathbf{T}(\tau)|. \quad (5)$$

The statistics of the environment are stored in $\mathbf{M}$ (see also Figure 1b). Each stimulus is encoded in the superposition of the states of temporal history in which it was encoded. Again following TCM, we can then use the current state of the temporal history as a probe to predict what stimulus will be presented at the next time step:

$$|\mathbf{p}(\tau)\rangle = \mathbf{M} |\mathbf{T}(\tau)\rangle, \quad (6)$$

where $\mathbf{p}$ is referred to as the prediction vector. We can combine Eqs. 5 and 6 to observe that at each time step, each stimulus is predicted to the extent that it overlaps with its encoding context.

It should be noted that the summation on the rhs of Eq. 6 is weighted by $g(\tau^*)$ the number density of nodes representing a particular value of τ^* . If N indexes the number of the cell, $g(\tau^*) \equiv \frac{dN}{d\tau^*}$. More explicitly,

$$|\mathbf{p}\rangle = \int_{-\infty}^0 \mathbf{M}_{\tau}^* g(\tau^*) |\mathbf{T}_{\tau}^*\rangle, \quad (7)$$

where the subscript notation $\mathbf{M}_{\tau}^*$ refers to the matrix within the three-tensor $\mathbf{M}$ corresponding to a specific value of τ^* . Similarly, $|\mathbf{T}_{\tau}^*\rangle$ refers to the vector within $\mathbf{T}$ corresponding to a particular value τ^* .

Predicting the future using backward replay. It would be extraordinarily useful if we were able to predict not only the immediately preceding moment, but also to generate a trajectory of future states leading from the present. There are many ways that one could conceivably implement a representation like $\mathbf{T}$. One of the major advantages of the approach we've developed in the previous program period is that it provides a concise and physical mechanism for calculating such trajectories into the future.

Note that we can write a discrete version of the differential equation encoding the Laplace transform (Eq. 1)

$$\mathbf{t}_{\tau+1}(s) = \mathbf{R} \mathbf{t}_{\tau}(s) + \mathbf{f}_{\tau}. \quad (8)$$

Here the operator $\mathbf{R}$ (in analogy with ρ in TCM) is just a diagonal matrix with $\rho \equiv e^{-s}$ on the row corresponding to each value of s in the sheet $\mathbf{t}(s)$.

In order to construct an estimate of what stimulus will be available δ steps in the future, we need to estimate the state of $\mathbf{T}$ that will obtain at that time. That is, we can estimate $\mathbf{p}(\tau + \delta)$ by operating on $\mathbf{T}(\tau + \delta)$ with $\mathbf{M}$. The difficulty is obtaining $\mathbf{T}(\tau + \delta)$ at time τ without waiting δ additional time steps. As it turns out, this can be accomplished by simply altering the weights in $\mathbf{L}_k^{-1}$. To see how this is possible note that

$$\mathbf{T}(\tau + \delta, \tau^*) = \mathbf{L}_k^{-1} \mathbf{t}(\tau + \delta, s) \quad (9)$$

$$= \mathbf{L}_k^{-1} [\mathbf{R}^{\delta} \mathbf{t}(\tau, s)] \quad (10)$$

$$= [\mathbf{L}_k^{-1} \mathbf{R}^{\delta}] \mathbf{t}(\tau, s) \quad (11)$$




Figure 2. Schematic illustrating important features necessary for predicting trajectories forward in time. **a.** Schematic shows $\mathbf{t}(s)$ for an idealized situation when the columns correspond to distinct stimuli in the order in which they were presented. Note the exponential gradient. The operator that approximates inversion of the Laplace transform, $\mathbf{L}_k^{-1}$ is multiplied by $\mathbf{R}^\delta$. As a result $\mathbf{T}$ is shifted to approximate where those components would be δ steps in the future. Compare to Figure 1a. **b.** Characterization of the change to the weights of $\mathbf{L}_k^{-1}$. The center represents the weight along the diagonal. Flanking values give the weights of the adjacent positions within a row. Dark curve: $\mathbf{L}_k^{-1}$. Lighter curve: $\mathbf{R}^\delta \mathbf{L}_k^{-1}$. When operated on by $\mathbf{R}^\delta$, the weights are shifted along the diagonal and reduced in magnitude (not to scale).

That is, we can understand $\mathbf{R}^\delta$ as operating on $\mathbf{t}(\tau, s)$ as in Eq. 10, or we can understand it as being absorbed into the weights of $\mathbf{L}_k^{-1}$ as in Eq. 11. That is, because $\mathbf{T}$ is constructed from $\mathbf{t}(s)$ using known weights, and because $\mathbf{t}(s)$ evolves deterministically, it is possible to infer the $\mathbf{T}$ that would obtain in the future and express it simply by transiently changing the weights of $\mathbf{L}_k^{-1}$ appropriately.

A principled form for $g(\tau^)$.* In the initial publications describing $\mathbf{T}$, we kept $g(\tau^*)$ in Eq. 7 as a general function. However, it turns out that there is a principled choice for $g(\tau^*)$. Consider the distribution across τ^* caused by a more recent stimulus and by a less recent stimulus. For the less recent stimulus, the function will be more broad about the peak. It does not make sense to devote as many cells to the broad peak as to the less broad peak. We did a calculation in which we determined the function $g(\tau^*)$ that equalizes the information per cell. This turns out to be $g(\tau^*) = \frac{1}{\tau^*}$. Interestingly, this implies that the cells are logarithmically distributed. That is, internal, retrospective time obeys the Weber-Fechner law. Simultaneously, timing going forward (as in, say, a temporal bisection task), obeys the scalar property, with a linear relationship between predicted time and physical time. This is possible because forward-going timing is governed by the match of an entire $\mathbf{T}$ summed over τ^* . Both the encoding $\mathbf{T}$ and the retrieval $\mathbf{T}$ are logarithmically distributed, so scalar timing is observed (see Howard et al., In preparation, for details). Subsequent work has shown that this choice for $g(\tau^*)$ also equalizes the mutual information between adjacent cells within $\mathbf{T}(\tau^*)$ (Shankar & Howard, submitted).

Behavioral applications

The formalism described above provides a framework for remembering a history leading up to the present moment and generating a future trajectory. Here we summarize some highlights of the behavioral applications of the model we developed in the preceding period.




Figure 3. Hacker (1980) gave subjects a rapidly-presented list of letters followed by a choice between two probe letters. Subjects were instructed to choose the more recent probe. Reaction times for correct responses, responses where the subject chooses the more recent probe, depend very much on the recency of the more recent probe but are essentially independent of the recency of the less recent probe. In contrast, error RTs, in which the subject selects the less recent probe, depend very much on the recency of the less recent probe but are essentially unaffected by the recency of the more recent probe. Top: data. Bottom: Model. This is a one-parameter fit of the model.

Episodic memory across scales. Having a representation of history enables a description of previously puzzling results. It also allows for an account of phenomena across scales that had previously been attributed to separate memory stores operating at different time scales. We have applied the model extensively to two episodic memory tasks, free recall and the judgment of recency task. In free recall, the model generates scale-free recency and contiguity (Howard et al., In preparation). While scale-free recency has been assumed by other models (e.g., Brown, Neath, & Chater, 2007), we provided a principled process model that generates it as an emergent property of the equations given above. Scale-invariant recency at the behavioral level is a natural consequence of scale-invariance at the level of the representation. Scale-persistent contiguity is accomplished in our framework by assuming that there can be a recovery of the $\mathbf{t}(s)$ that obtained at a previous moment in time. This in turn enables recovery of the state of $\mathbf{T}$ that obtained at that time. Unlike previous attempts using TCM which have a scale fixed by the specific value of ρ used, this contiguity effect persists across arbitrarily large scales. It is not precisely scale-invariant because the asymmetry observed in the contiguity effect decreases as the scale is increased. This appears to be consistent with experimental data using final free recall (Unsworth, 2008).

The major difference between $\mathbf{T}(\tau^*)$ and previous models of temporal context is the fact that




Figure 4. Yntema and Trask (1963) gave subjects a continuous relative judgment of recency task. The figure shows the probability of choosing a probe item as more recent as a function of its recency and the recency of the other probe item. Model predictions use the same scanning model as used in Figure 3. The model has one free parameter.

$\mathbf{T}^*(\tau)$ retains temporal information about *when* the preceding stimuli were presented. In order to illustrate the advantage of this approach, we decided to model the judgment of recency (JoR) task. In the relative JoR task, subjects are given two probe stimuli and asked to choose the one that was presented more recently. With fast presentation of short lists, the JoR task generates results that have generally been taken as evidence for serial scanning. The data appear as if the subject examines the historical record in their memory moment-by-moment, stopping when they find either of the probes (Figure 3, Hacker, 1980; McElree & Doshier, 1993; Muter, 1979). We were able to account for these findings (Howard et al., In preparation) simply by assuming that the subject examines $\mathbf{T}^*(\tau)$ at successive values of τ^* starting at zero and extending backwards in time. At each moment, the subject detects a memory for a probe stimulus with probability proportional to the value of $\mathbf{T}$ in the appropriate stimulus column. When the subject detects a memory for one of the probes, she chooses it; if the scan goes on for long enough without detecting a memory, the subject guesses. This extremely simple model is able to account for the qualitative pattern of results in accuracy, mean correct RT and error RT (Figure 3).

Most scholars of memory have treated the Hacker (1980) results as evidence for a dedicated short-term store that is contrasted with long-term memory. However, we showed that the simple model described above also accounts for several phenomena from what are usually thought of as long-term JoR tasks. Figure 4 shows that the same model provides an excellent qualitative description of the experimental results of Yntema and Trask (1963), who performed a continuous relative JoR task with recencies chosen from a very broad temporal range.

In addition, we provided a principled account of several other findings in what are usually thought of as long term JoR tasks, including the logarithmic increase in absolute JoR's, in which subjects provide a numerical estimate for how many steps in the past a stimulus was presented (Hinrichs & Buschke, 1968; Hinrichs, 1970) and the approximate independence of separate judgments based on distinct presentations of an item (Hintzman, 2010). The former follows from the choice of $g(\tau^*)$ and the latter follows from the linearity of $\mathbf{L}_k^{-1}$. We also showed that failures to observe substantial performance in the relative JoR task (Klein, Shiffrin, & Criss, 2007) can be accounted for by the choice of relative and absolute delays those authors used.



Figure 5. Procedure a typical temporal mapping experiment. **a.** In the first phase of training, one group (top) received CS1 with the shock US presented immediately after CS1 offset. The other group (bottom) received the US five seconds after offset of CS1. **b.** Both groups received second-order conditioning between CS1 and CS2. **c.** The presumed representation after both phases of training for each group if the experiences were aligned on CS1 onset. In the second group, CS2 predicts US onset more robustly than in the first group. After Cole, Barnet & Miller (1995).

Temporal mapping. Ralph Miller and colleagues have done a series of elegant conditioning experiments that suggest two fundamental properties of learning that are particularly well-suited for the current mathematical framework. First, they argue that the temporal relations between stimuli forms an essential and unavoidable part of the learning event. Second, they argue that learners can integrate disparate learning events into a coherent temporal map by aligning different time lines on a common stimulus.

To make this more concrete, let us describe a specific experiment (Cole et al., 1995). Rats were trained to associate a 5 s CS1 with a US (shock). In one condition, the time between offset of the CS1 and the onset of the US was 0 s. In the other condition, the time between the offset of CS1 and the US was 5 s (Figure 5a). Let us refer to these as the 0 s and 5 s conditions, respectively. After training the CS1-US association, a second-order association was formed between CS1 and another 5 s CS2. In both conditions, the onset of CS2 immediately followed the offset of CS1 (Figure 5b). In neither condition did CS2 ever cooccur with the US. The first finding was, not surprisingly, that the CR to the CS1 was stronger in the 0 s condition than in the 5 s condition. If associative strength is a scalar value, we would expect the second order conditioning to CS2 would also be stronger in the 0 s condition than in the 5 s condition. However, exactly the opposite was observed. This result makes no sense from the perspective of simple associative strength. Miller's temporal coding hypothesis (Matzel, Held, & Miller, 1988; Savastano & Miller, 1998) reconciles these findings as follows. Note that if the two learning episodes were aligned on the CS1 (as in Figure 5c), then the CS2 does not predict the onset of the US in the 0 s condition. In the 5 s condition, CS2 strongly predicts the onset of the US when the two learning episodes are aligned.

A model must have two basic properties in order to account for this phenomenon. One is that the temporal relationships between stimuli, rather than a simple scalar associative strength, is learned. Second, some mechanism for integrating disparate episodes into a coherent synthetic representation is necessary. The representation of temporal history offered by big **T** satisfies the first constraint. The ability to retrieve temporal contexts satisfies the second constraint (see also Howard et al., 2005; Rao & Howard, 2008). We showed that the mathematical properties of the model enable it to construct satisfactory temporal maps between disparate experiences (Howard et al., In preparation).

References

- Brown, G. D., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, 114(3), 539-76.
- Cole, R. P., Barnet, R. C., & Miller, R. R. (1995). Temporal encoding in trace conditioning. *Animal Learning & Behavior*, 23(2), 144-153.
- Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J., & Usher, M. (2005). The demise of short-term memory revisited: empirical and computational investigations of recency effects. *Psychological Review*, 112(1), 3-42.
- Farrell, S. (2012). Temporal clustering and sequencing in short-term memory and episodic memory. *Psychological Review*, 119(2), 223-71.
- Glenberg, A. M., Bradley, M. M., Stevenson, J. A., Kraus, T. A., Tkachuk, M. J., & Gretz, A. L. (1980). A two-process account of long-term serial position effects. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 355-369.
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114(2), 211-44.
- Grossberg, S., & Pearson, L. R. (2008). Laminar cortical dynamics of cognitive and motor working memory, sequence learning and performance: toward a unified theory of how the cerebral cortex works. *Psychological Review*, 115(3), 677-732.
- Hacker, M. J. (1980). Speed and accuracy of recency judgments for events in short-term memory. *Journal of Experimental Psychology: Human Learning and Memory*, 15, 846-858.
- Hinrichs, J. V. (1970). A two-process memory-strength theory for judgment of recency. *Psychological Review*, 77(3), 223-233.
- Hinrichs, J. V., & Buschke, H. (1968). Judgment of recency under steady-state conditions. *Journal of Experimental Psychology*, 78(4), 574-579.
- Hintzman, D. L. (2010). How does repetition affect memory? Evidence from judgments of recency. *Memory & Cognition*, 38(1), 102-15.
- Howard, M. W., Fotedar, M. S., Datey, A. V., & Hasselmo, M. E. (2005). The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychological Review*, 112(1), 75-116.
- Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 923-941.
- Howard, M. W., & Kahana, M. J. (2002). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, 46(1), 85-98.
- Howard, M. W., Shankar, K. H., Aue, W., & Criss, A. H. (In preparation). A quantitative model of time in episodic memory.
- Howard, M. W., Shankar, K. H., & Jagadisan, U. K. K. (2011). Constructing semantic representations from a gradually-changing representation of temporal context. *Topics in Cognitive Science*, 3, 48-73.
- Howard, M. W., Viskontas, I. V., Shankar, K. H., & Fried, I. (In press). A neural signature of mental time travel in the human MTL. *Hippocampus*.
- Howard, M. W., Youker, T. E., & Venkatadass, V. (2008). The persistence of memory: Contiguity effects across several minutes. *Psychonomic Bulletin & Review*, 15, 58-63.
- Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, 24, 103-109.
- Kılıç, A., Criss, A. H., & Howard, M. W. (In press). A causal contiguity effect that persists across time scales. *Journal of Experimental Psychology: Learning, Memory and Cognition*, <http://dx.doi.org/10.1037/a0028463>.
- Klein, K. A., Shiffrin, R. M., & Criss, A. H. (2007). Putting context in context. In J. S. Nairne (Ed.), *The foundations of remembering: Essays in honor of Henry L. Roediger, III*. (p. 171-189). New York, NY: Psychology Press.
- Landauer, T. K., & Dumais, S. T. (1997). Solution to Plato's problem : The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211-240.

- Matzel, L. D., Held, F. P., & Miller, R. R. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning and Motivation*, 19, 317-344.
- McElree, B., & Doshier, B. A. (1993). Serial recovery processes in the recovery of order information. *Journal of Experimental Psychology: General*, 122, 291-315.
- Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, 64, 482-488.
- Muter, P. (1979). Response latencies in discriminations of recency. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 160-169.
- Post, E. (1930). Generalized differentiation. *Transactions of the American Mathematical Society*, 32, 723-781.
- Rao, V. A., & Howard, M. W. (2008). Retrieved context and the discovery of semantic structure. In J. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in neural information processing systems 20* (p. 1193-1200). Cambridge, MA: MIT Press.
- Savastano, H. I., & Miller, R. R. (1998). Time as content in Pavlovian conditioning. *Behavioural Processes*, 44, 147-162.
- Sederberg, P. B., Howard, M. W., & Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review*, 115, 893-912.
- Shankar, K. H., & Howard, M. W. (2012). A scale-invariant representation of time. *Neural Computation*, 24, 134-193.
- Shankar, K. H., & Howard, M. W. (submitted). Optimally fuzzy scale-free memory. In *Advances in neural information processing systems*.
- Shankar, K. H., Jagadisan, U. K. K., & Howard, M. W. (2009). Sequential learning using temporal context. *Journal of Mathematical Psychology*, 53, 474-485.
- Unsworth, N. (2008). Exploring the retrieval dynamics of delayed and final free recall: Further evidence for temporal-contextual search. *Journal of Memory and Language*, 59, 223-236.
- Yntema, D. B., & Trask, F. P. (1963). Recall as a search process. *Journal of Verbal Learning and Verbal Behavior*, 2, 65-74.